ALCF

OLCF Oak Ridge Leadership Computing Facility

NERSC

SC2017, November 12-17, Denver
Session: 2nd International Workshop on Data Reduction on Big Scientific Data

Facing the big data challenge in the fusion code XGC

C-S Chang

Princeton Plasma Physics Laboratory
Princeton University

on behalf of
The SciDAC-4 Partnership Center for
High Fidelity Boundary Plasma Simulation

PPPL

XBP
High-fidelity
Boundary Plasma Simulation

ITER construction is well-underway

- Plasma: Ionized state of matter
- $D^{+1} + T^{+1}$ (>100M degrees) = α^{+2}(3.5MeV) + n(14MeV)

- Fuel: D_2 extracted from body of water
  - Enough for millions of years
- Tritium is self-bred
- Melt-down cannot happen
  - If anything goes wrong, we lose the fusion reaction
  - Source of difficulty
- Wall-activation by neutrons: decays to safe level <100yrs
  - Unlike ~10,000 yrs for fission reactor waste

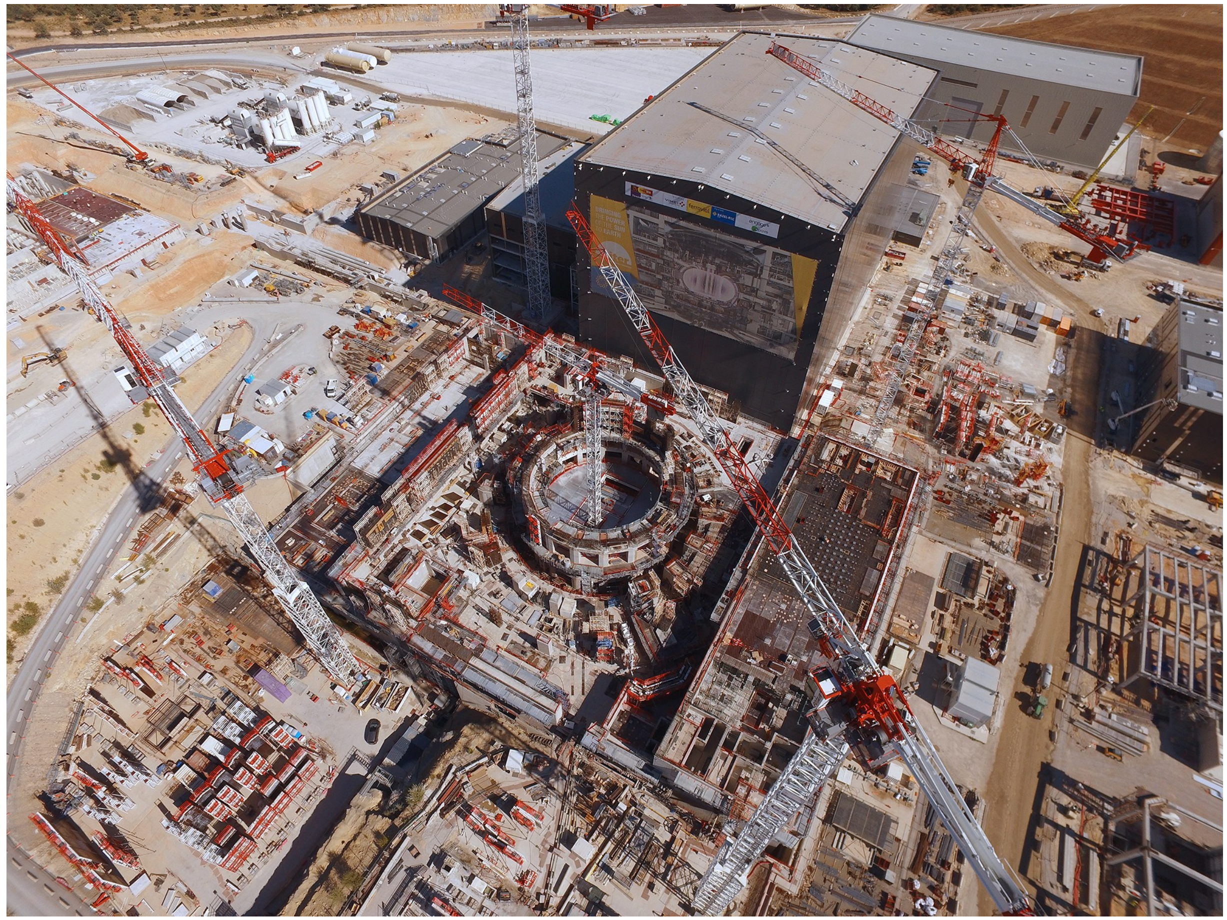

Outline

- Plasma confinement physics in tokamak
- Why does an edge plasma simulation in XGC produce an uncontrollable amount of big data?
- Examples on what we are looking for in the sea of data.
- A new paradigm is needed for analyzing and validating the XGC simulation data

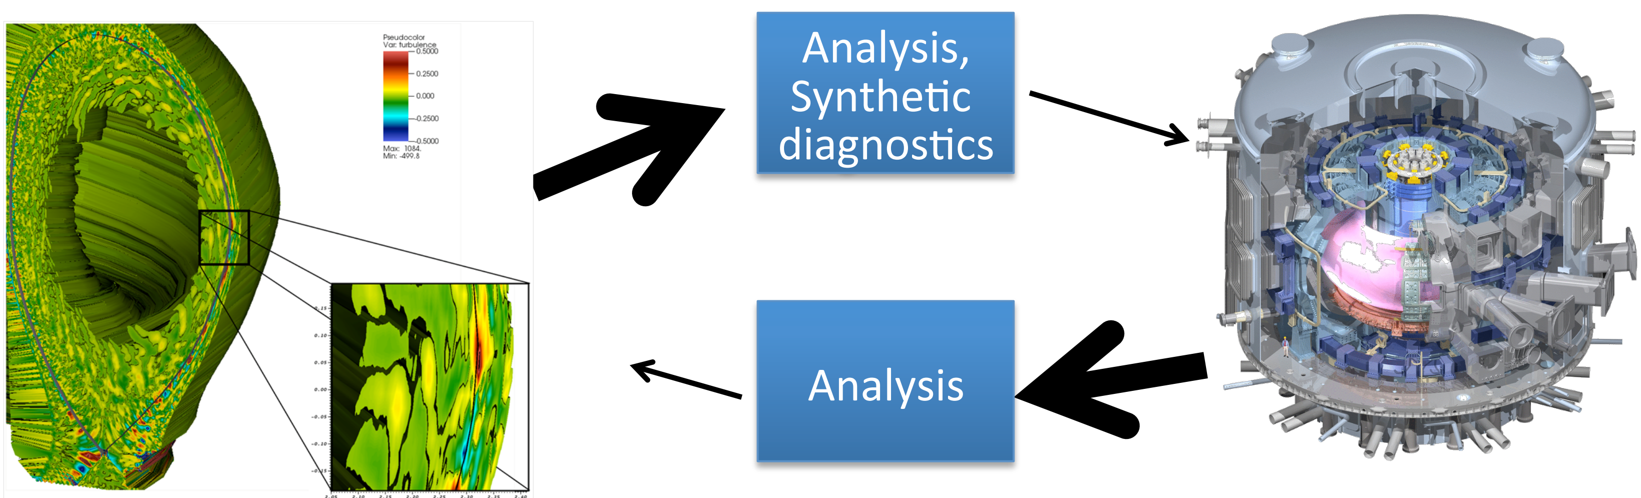

ITER

Simulation by S. Ku, PPPL
Visualization by D. Pugmire, ORNL

“Toroidal” Tokamak Geometry

Poloidal magnetic flux label
$\psi(r) = 1$ at $r/a=1$, 0 at $r/a=0$

ITER

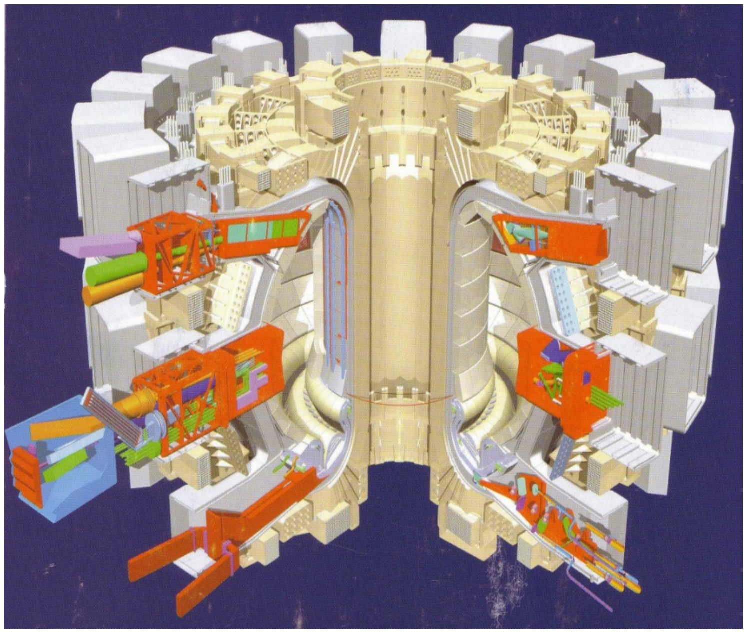

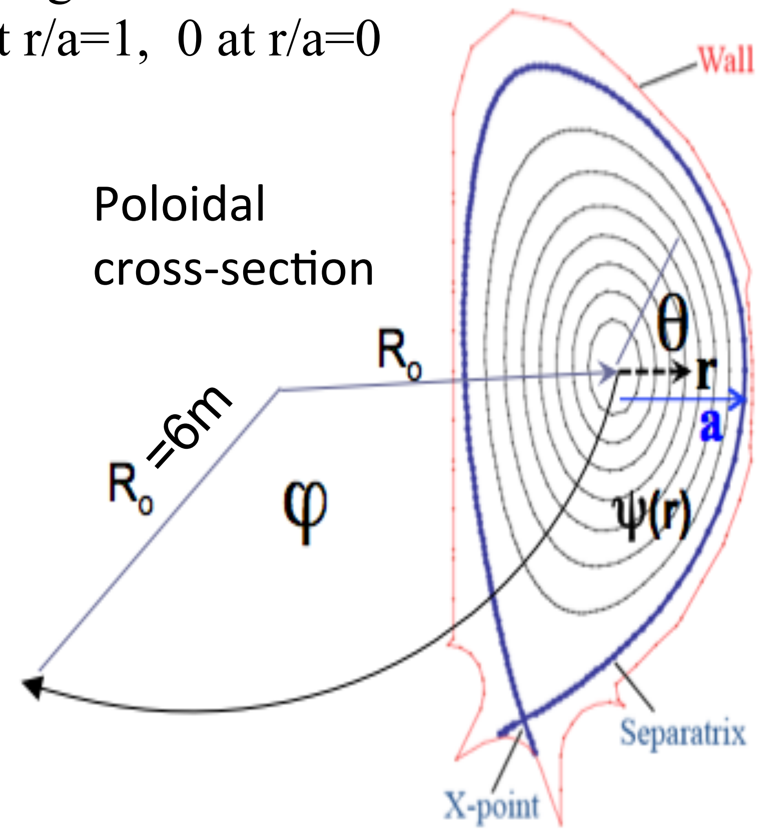

Torus, not a straight cylinder: physics and math become more challenging: toroidal-poloidal mode coupling, non-local interactions, , ballooning, complicated particle motions......

Near-thermal-equilibrium state vs far-from-equilibrium state

Closed system $\delta Q / Q \rightarrow 0$

δQ

$f \cong f_M + \delta f$

Q=energy in the system

δQ

$\delta Q << Q$

$\delta f << f_M$

- Could study δf ($\sim 0.01\ f_M$) only.
- Small and slow δf
- Small data and cheap simulation

Open system

δQ

$f \neq f_M$

Q

δQ

$\delta Q \sim Q$

$\delta f \sim f_M$

- Must study the whole f.
- Large and fast δf
- Big data and expensive simulation

In a homogeneous B-field, charged particles travel along B-line, while gyrating: Poincare's 3D torus → tokamak. In tokamak, however, plasma particles drift up or down.

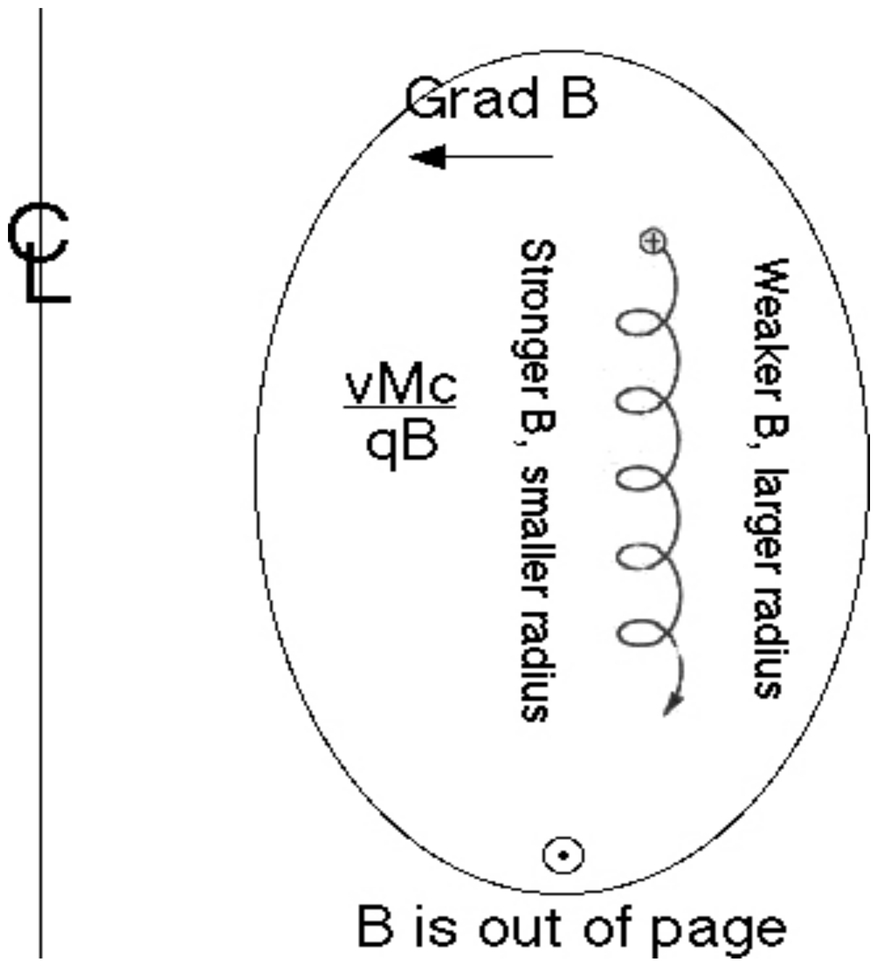

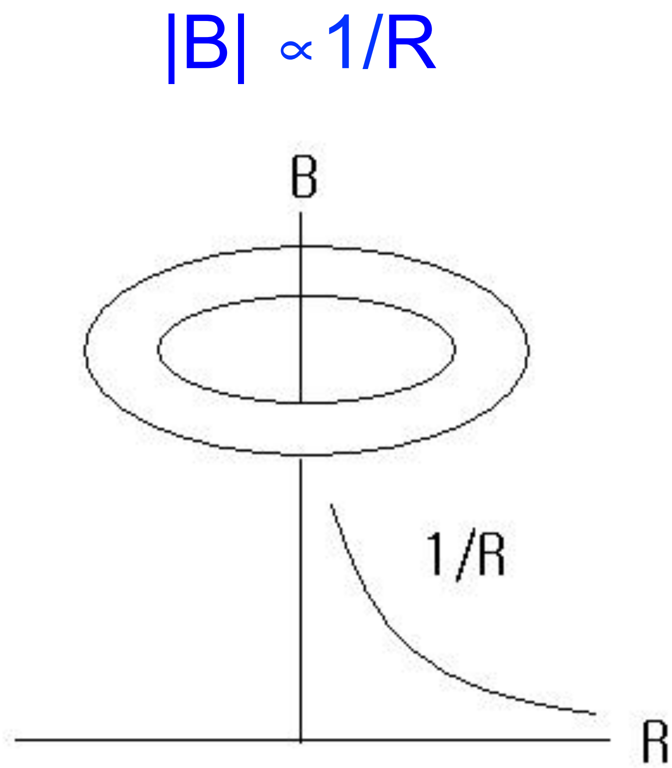

With the toroidal magnetic field only, plasma particles are not confined

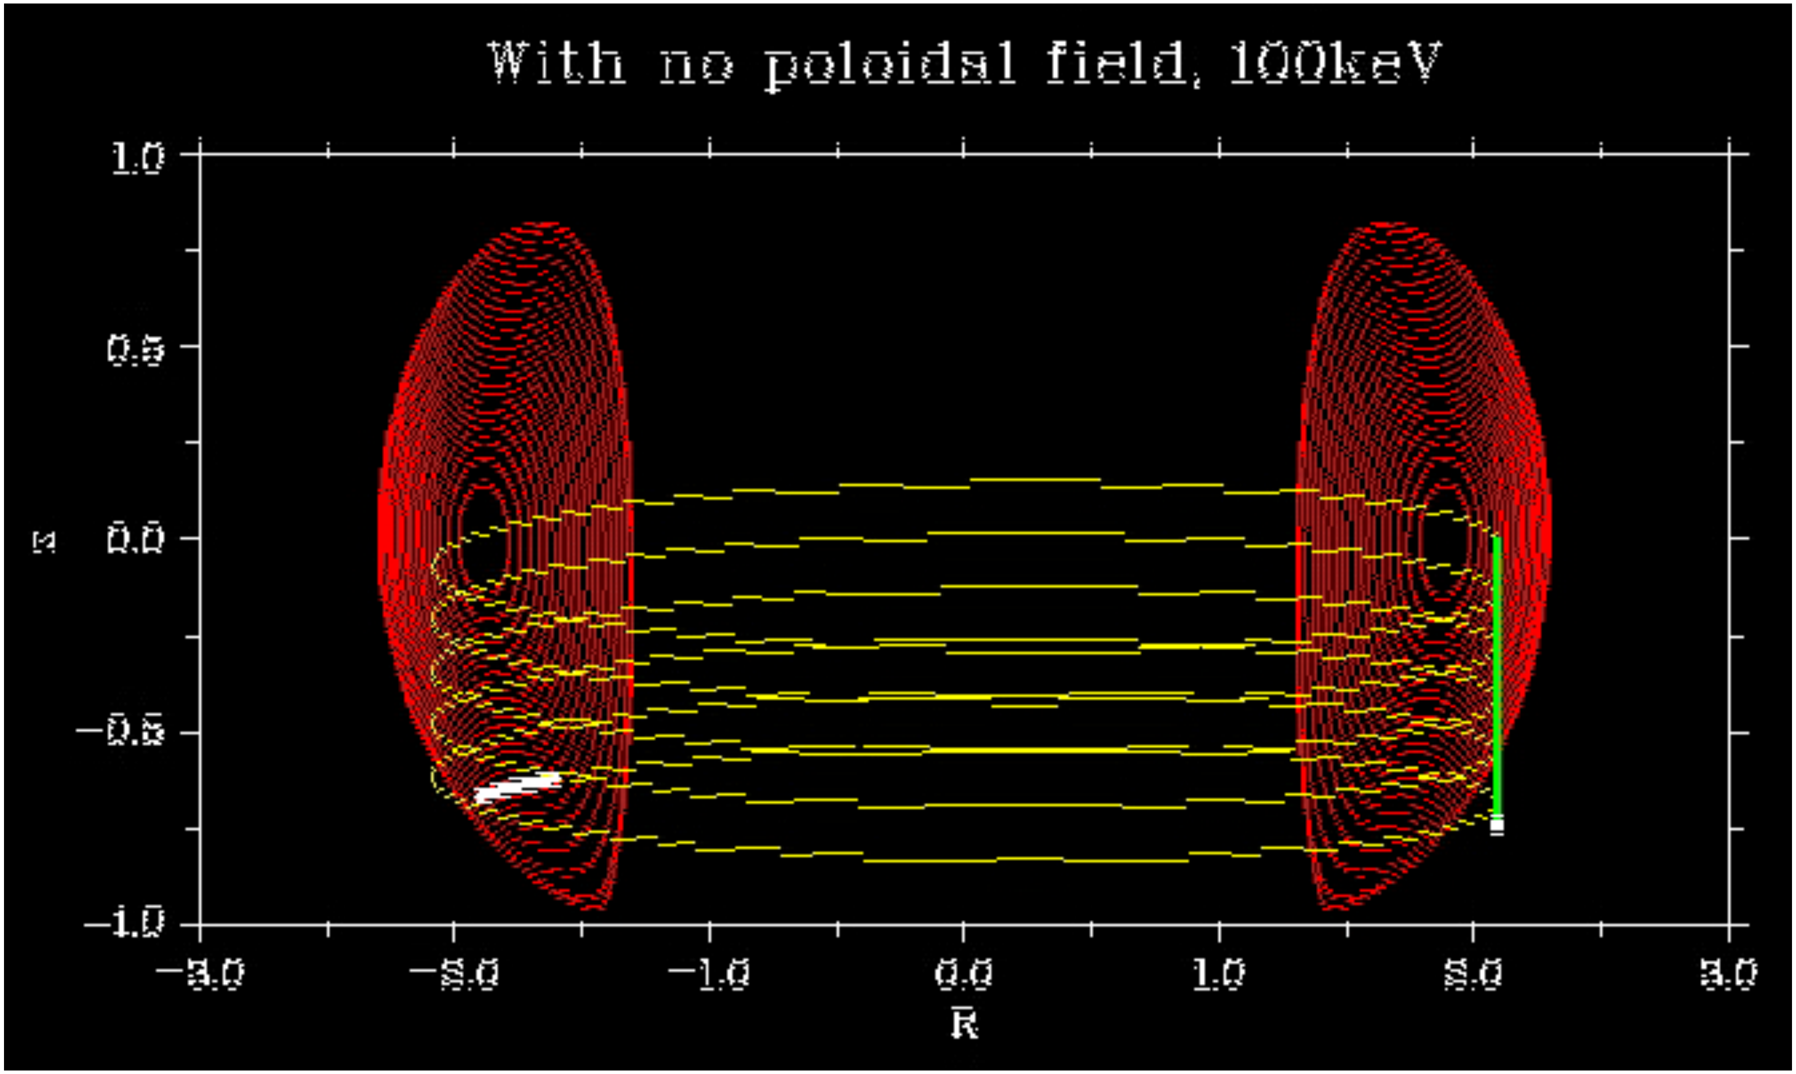

Gyro motions are not shown.

After adding poloidal magnetic field, plasma particles are confined.

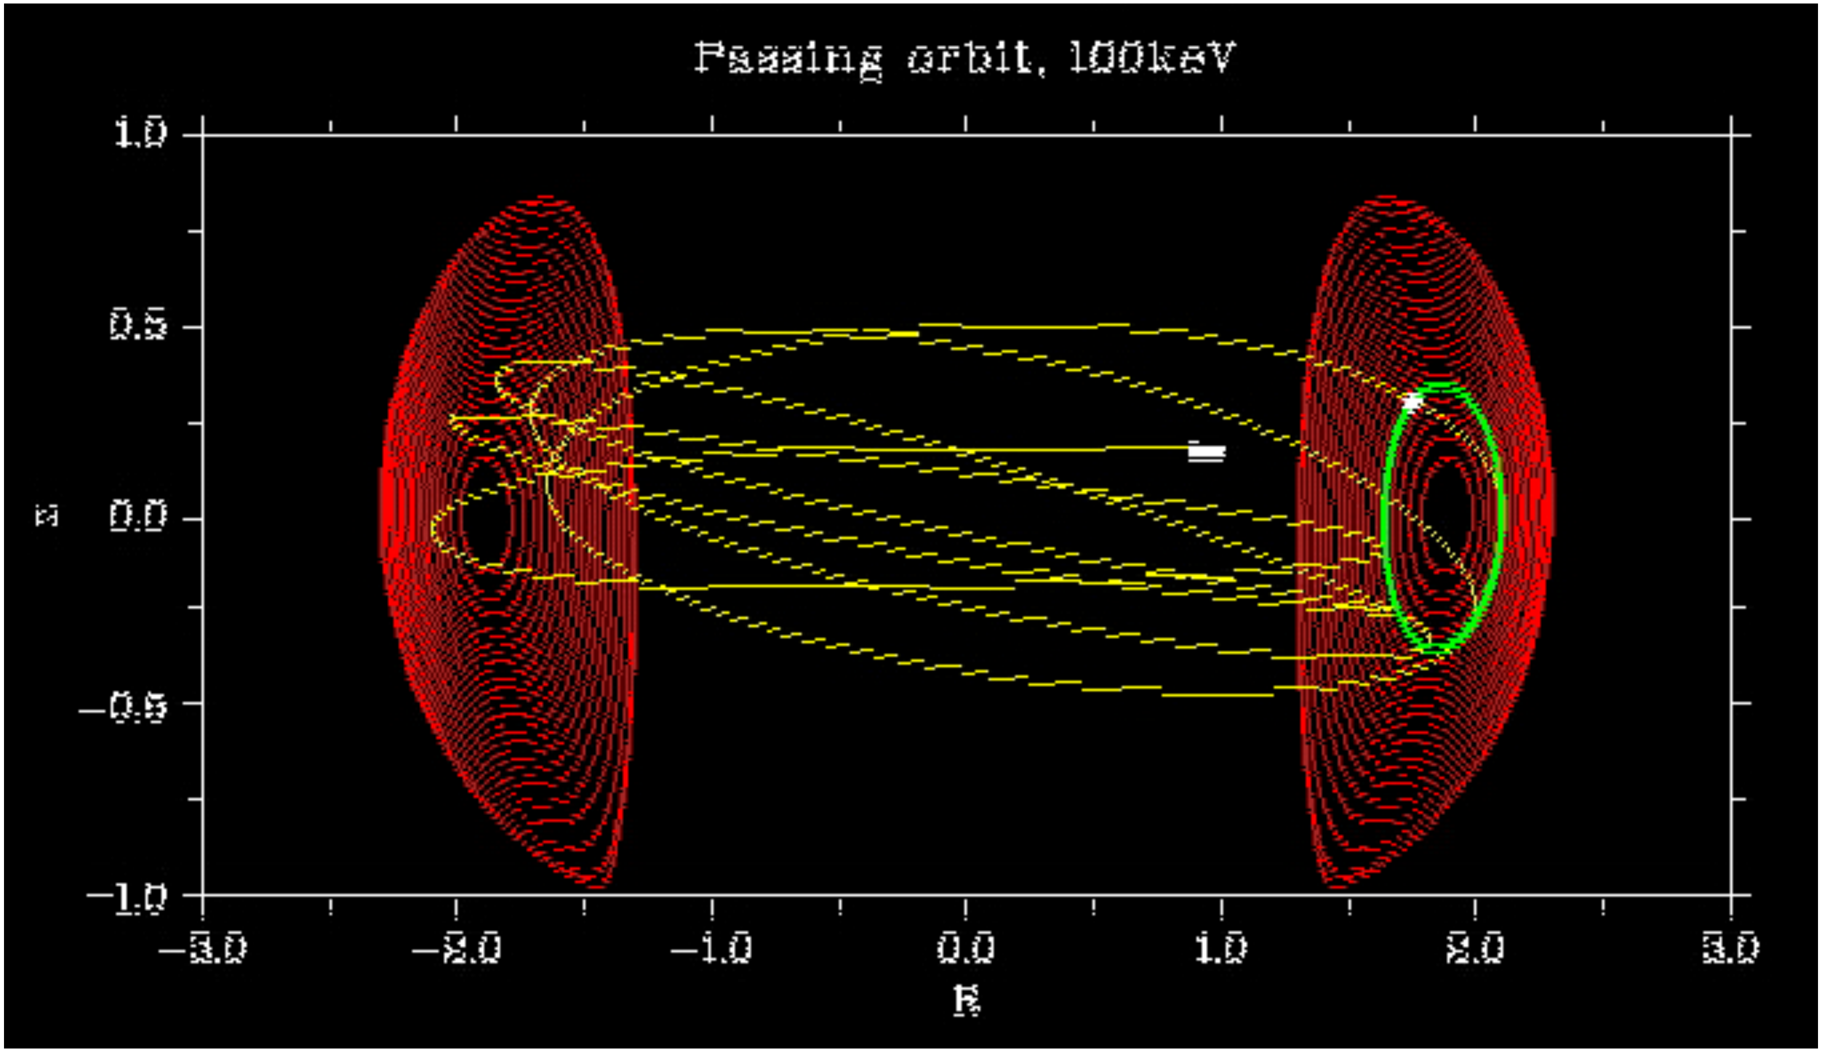

Gyro motions are not shown.

Magnetic mirror force turns particles with small $v_{||}$ into trapped "banana" orbits, which creates various nonlocal physics and/or different wave-particle interaction

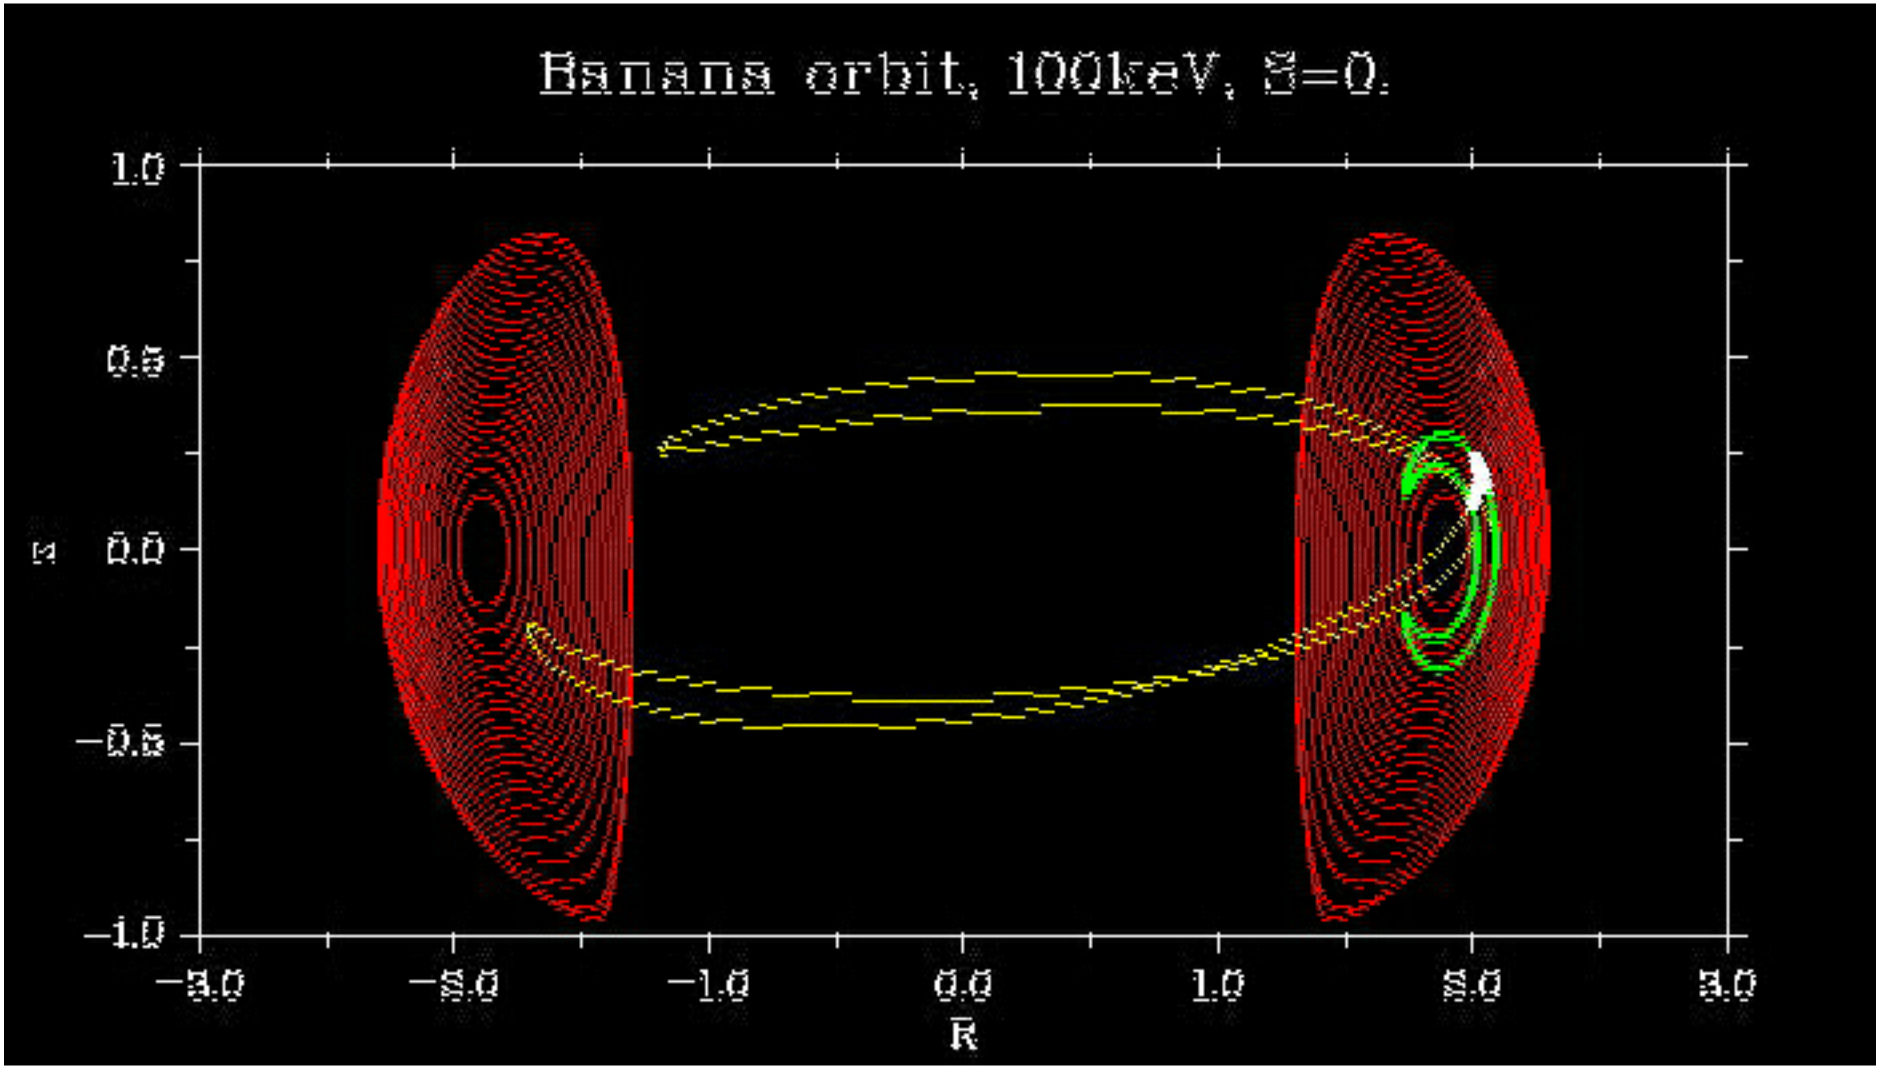

Gyro motions are not shown.

In the well-confined core, the plasma transports out its heat by the self-organized "streamer" activities

Ion Temperature Gradient turbulence

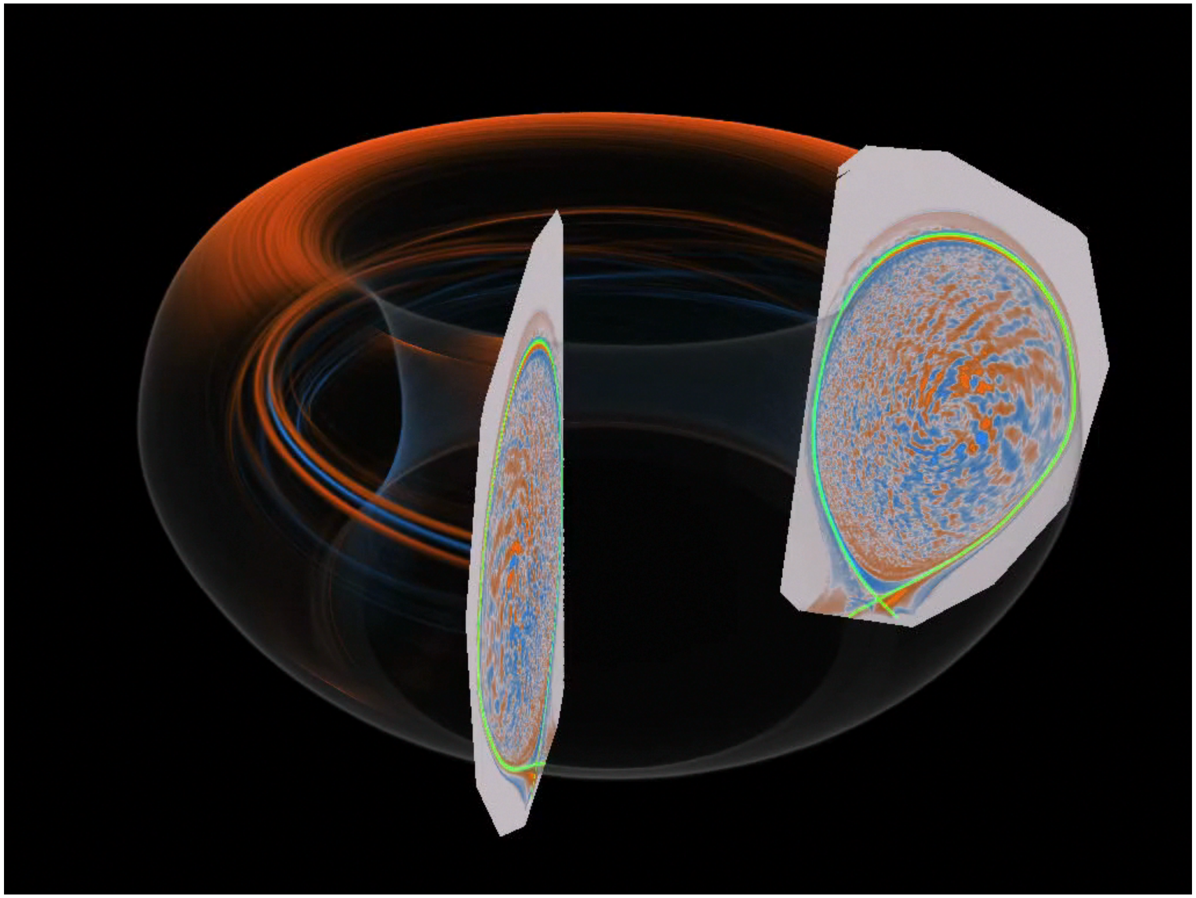

Global self-organization is established in core turbulence time ~ms

XGC1 simulation by S. Ku, Visualization by K. Ma

How does the streamers transport plasma?

-through the ExB-drift and dissipation-

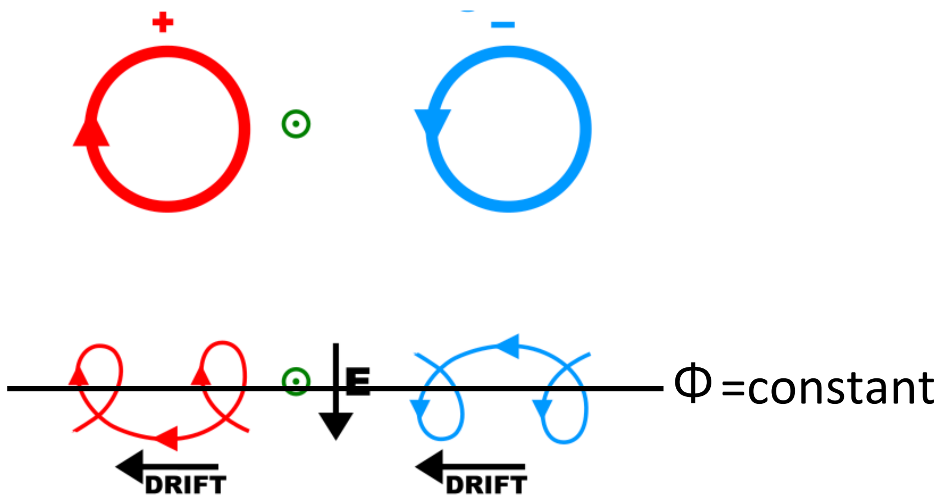

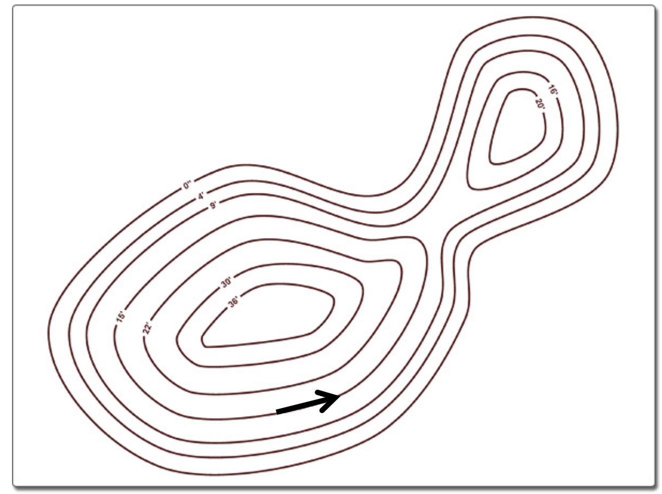

- Ions and electrons move along the constant electric potential contour lines
- Collisions and turbulence decorrelation provide “leaks” along the motion
- Radial transport rate is self-controlled via the turbulence and mean-ExB flow self-organization

Regulation of turbulent “streamer” size by sheared mean ExB-flow

There are two types of ExB flows

- The fluctuating one that generates “streamers” and yields radial transport
- The slow time-varying one that regulates the “streamers”

Understanding the dynamical interaction of these two ExB flows with turbulence is an important part of data analysis.

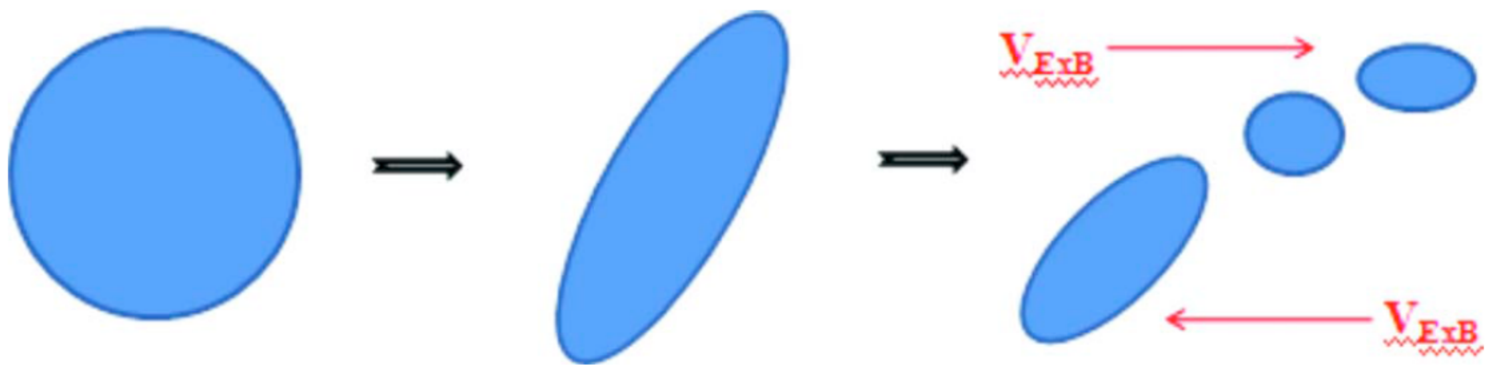

FIG. 1. (Color online) Cartoon illustration of a large scale eddy being sheared into small scale turbulence (see Ref. 18).

Self-regulation of turbulence and heat flux by sheared ExB flow

Can you identify the self-regulating activities?

L_{turb} cascades down to 2π ion-gyroradius ~ 1 cm

Zonal flows at the beginning of steady turbulence

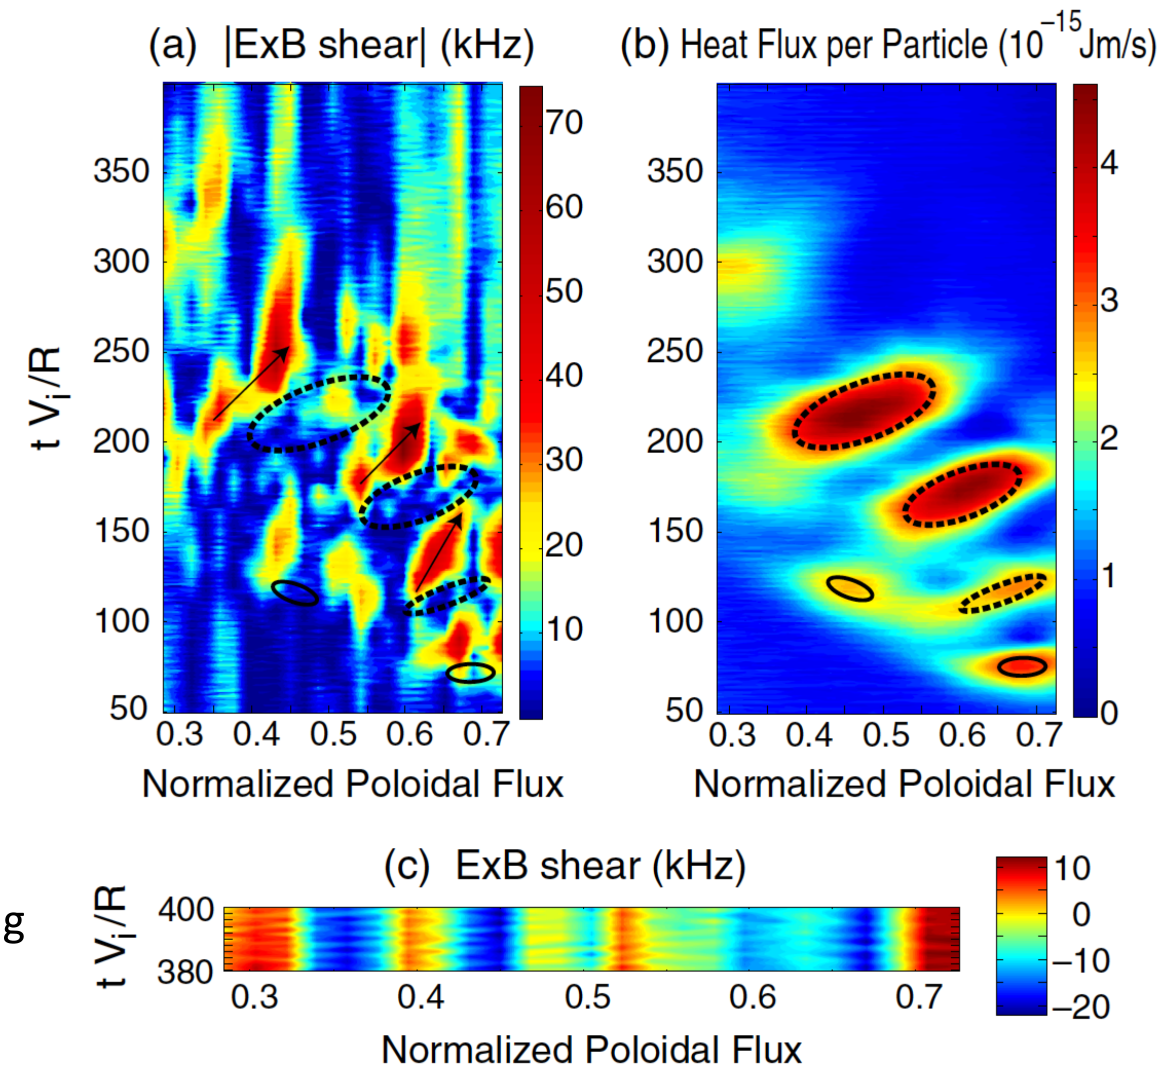

Example for post-processing analysis of coarse-grained data

Hidden Physics (on mesh)

- Self-organization is regulated by spontaneous ExB flow shearing, through non-local staircase T or corrugated ∇T profile interactions
- Sheared ExB flow and ∇T corrugation keeps the turbulence to be just right for to expel a proper amonut of heat-flux
- Non-Gaussian turbulence

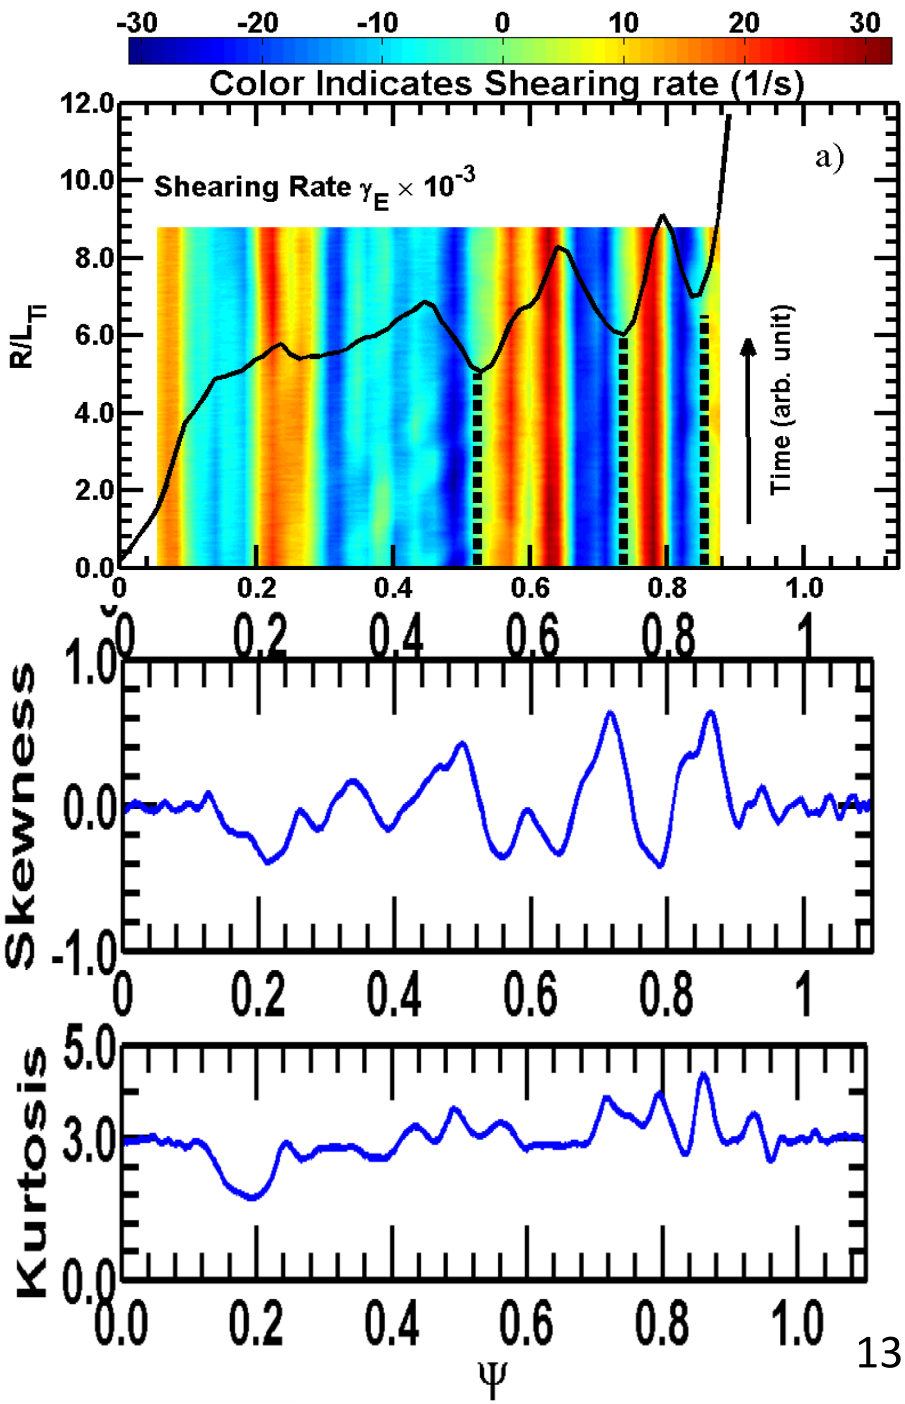

Success of ITER depends on this self-organization physics, built on top of the H-mode edge pedestal

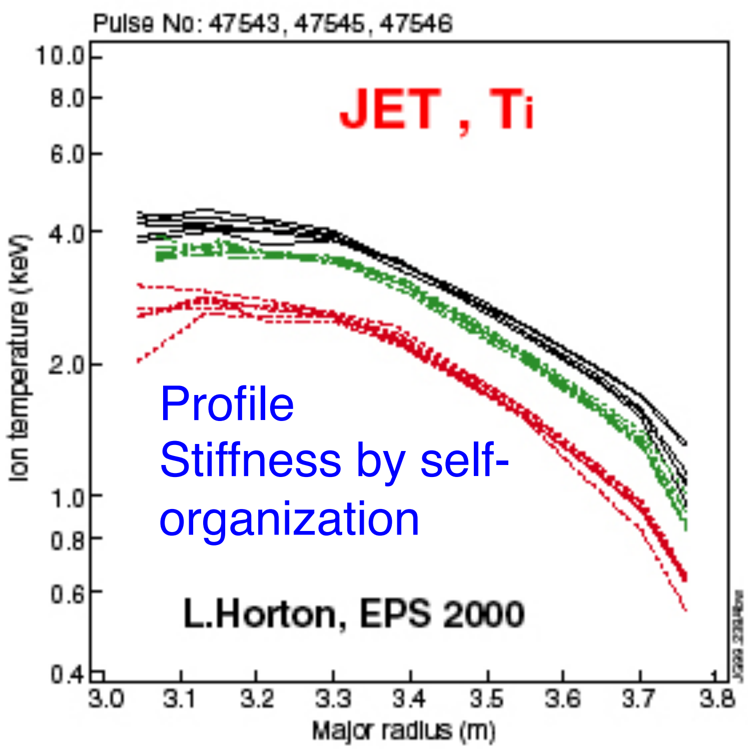

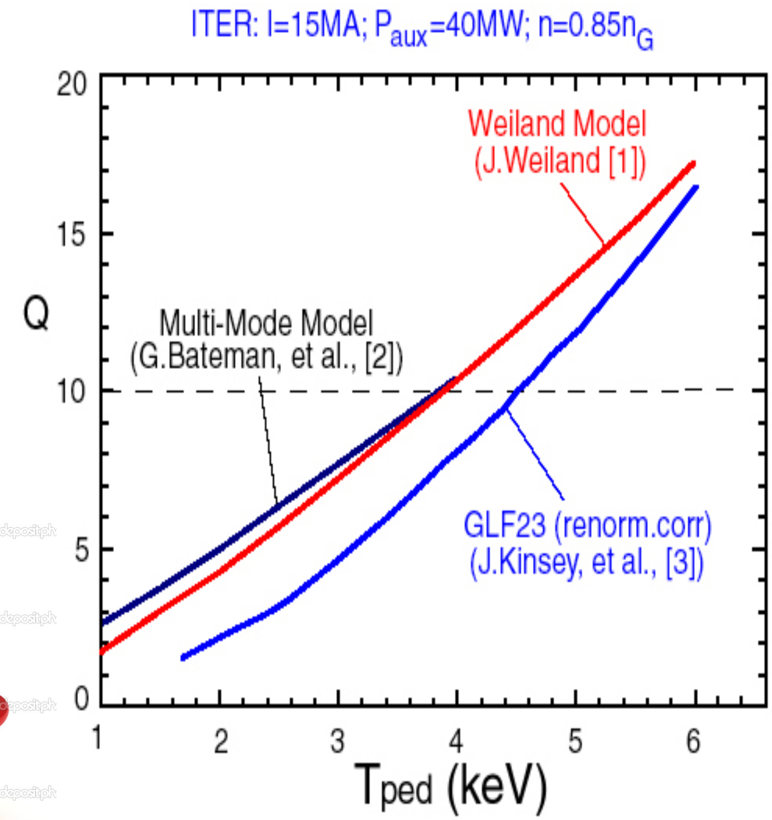

H-mode bifurcation is another mysterious self-organization process.

[C.S. Chang, PRL 2017]

As we approach the edge, we begin to lose particles even from inside the magnetic separatrix surface

- B_P=0 at magnetic X-point and is small around it.
  - Confinement is lost → X-point ion orbit loss
  - Negative charge within ion orbit width Δ_b inside separatrix → strong E_r<0 in Δ_b layer
- Strong V_{ErxB} restores poloidal rotation and restores the ion confinement
  - Stops further build-up of E_r
- Strong E_r yields turbulence bifurcation

  → H-mode pedestal

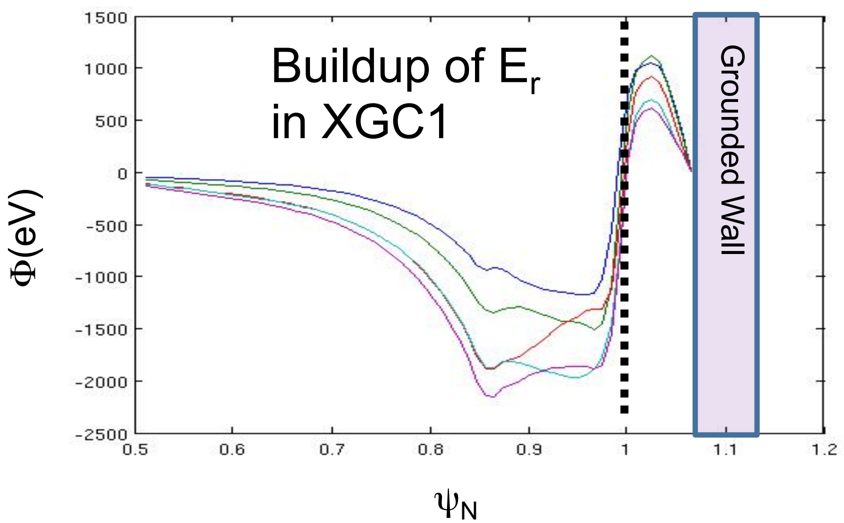

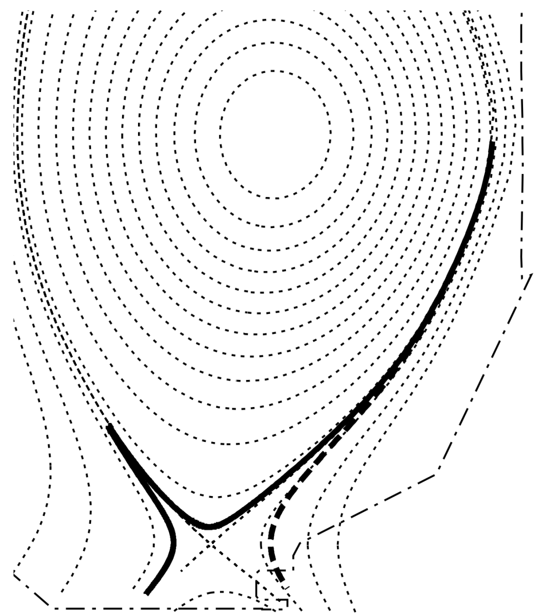

Typical X-point ion loss orbits, from XGC

Outside the magnetic separatrix surface, the B-field lines are completely open (purple lines) → far-from-equilibrium

Heat from core, particle source from the recycled neutral-particle ionization, and heat and particle loss to the wall.

→ Sharp radial gradient
→ All-scale physics must be solved together using the fundamental kinetic equation.
→ XGC kinetic code using trillion particles (ITER)
→ Big data

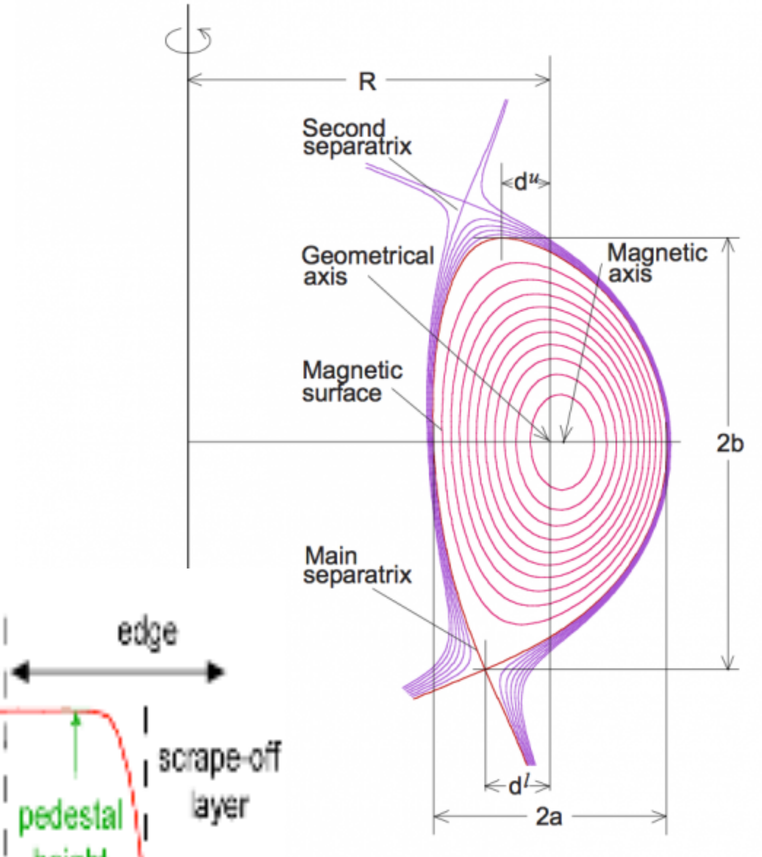

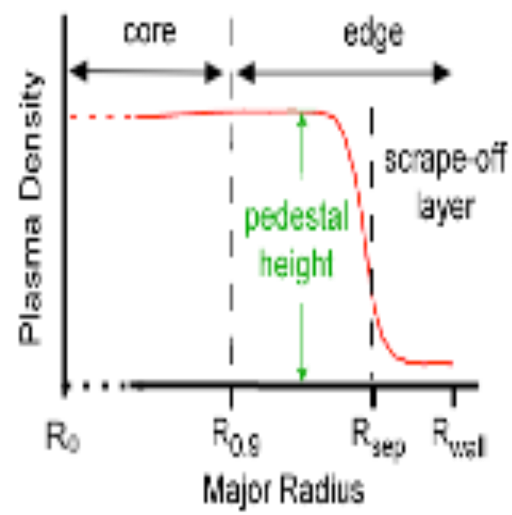

Edge Pedestal

XGC is an extreme scale code running on both LCF architectures at full capability.

- is in all three pre-exascale/ exascale programs (NESAP, CAAR, and Aurora ESP)
- currently burns ~500Mhrs per year for fundamental and critical physics research that cannot be performed by other codes

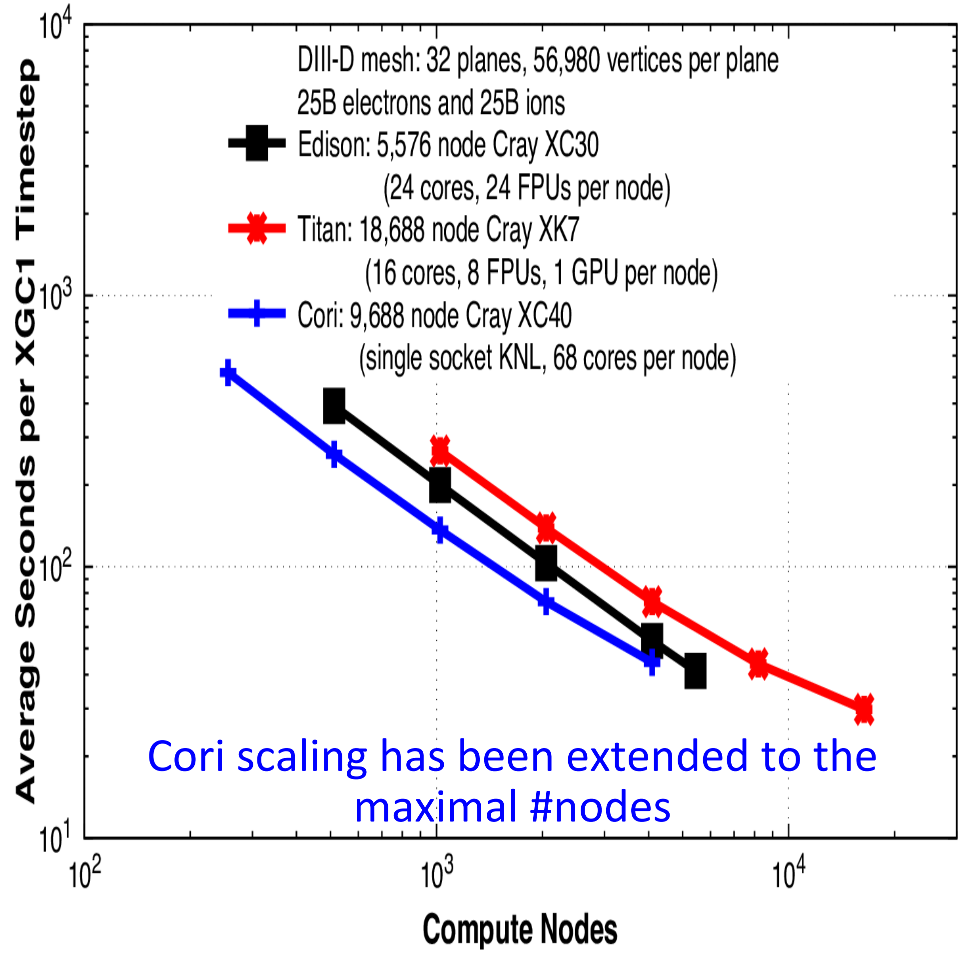

2D neutral particles evolve consistenly with plasma

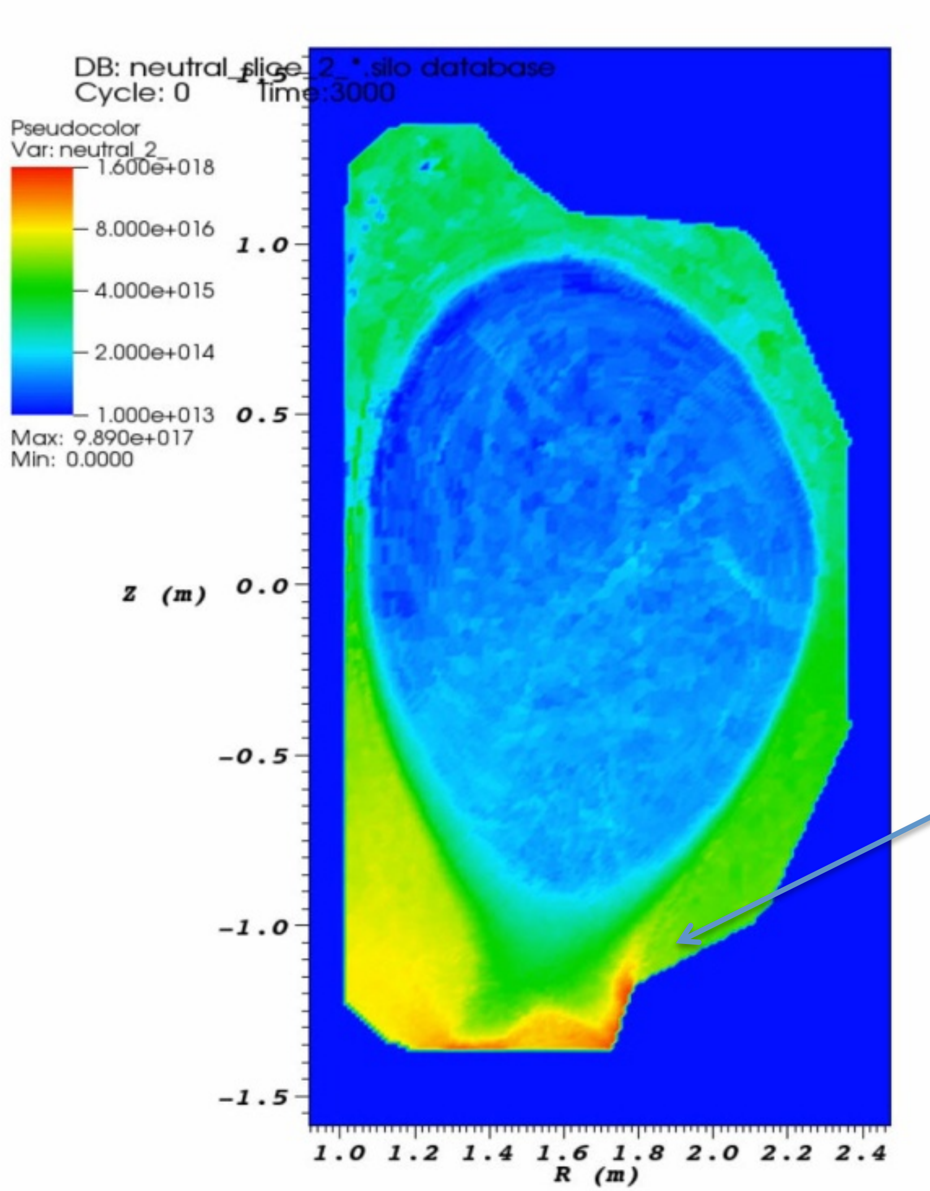

Figure by D. Stotler

Logarithmic plot of 2D deuterium neutral atom density in a DIII-D plasma

(showing that that the neutral source is peaked at the divertor targets, as determined by the poloidal profile of XGC ion losses to wall).

Edge tubulence activities contain a lot of important physics information that are critical to the success of ITER and the fusion reactors

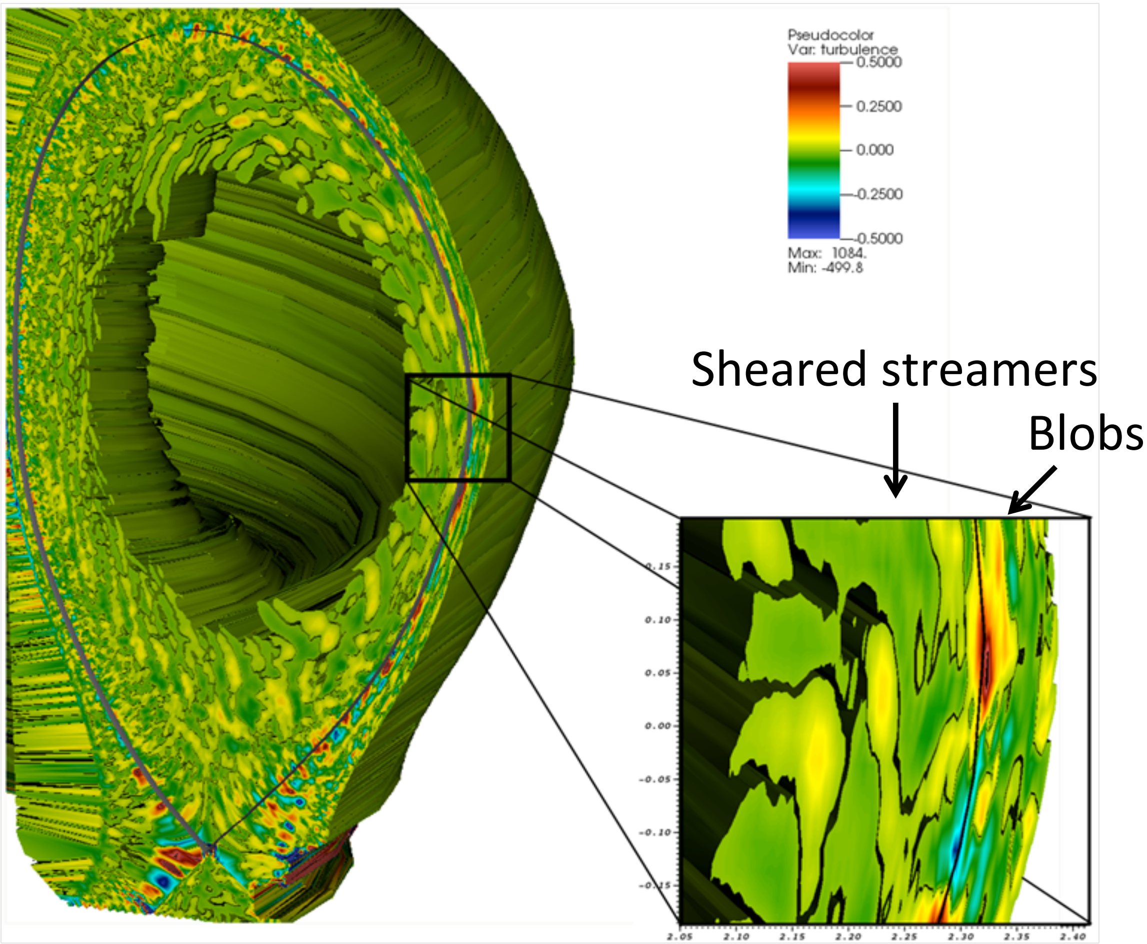

Blob dynamics in DIII-D like H-mode edge

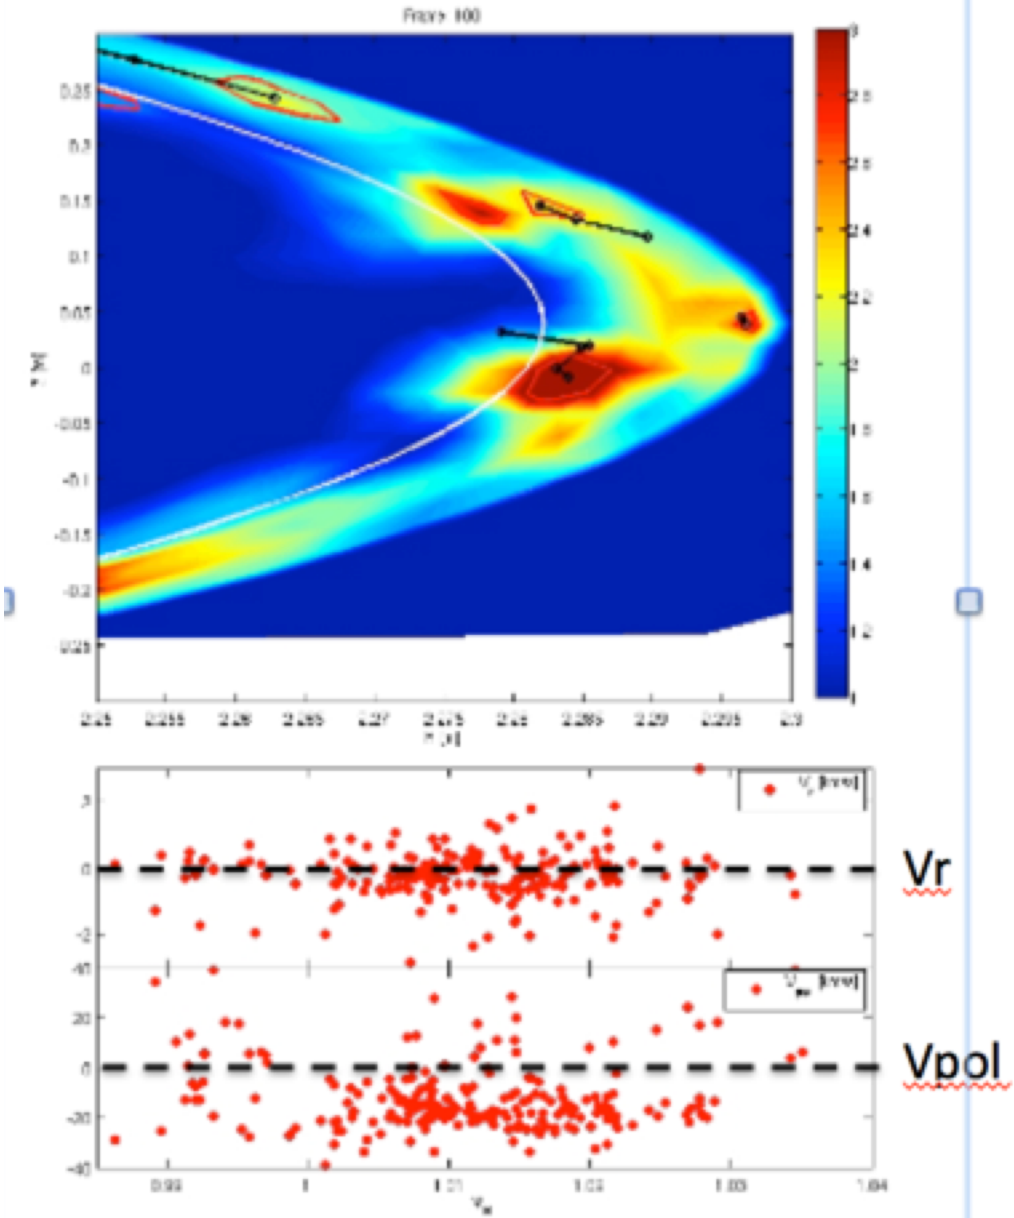

Blob radial velocity stays below 2 km/s.

Poloidal ExB blob motion is in the electron diamagnetic direction (upward) in the pedestal, and changes sign in the scrape-off layer toward the divertor (~20 km/s).

Why $\langle V_r \rangle \sim 0$? Baseline physics study of collisionless blob structure/dynamics [Churchill, PPCF, submitted]: Blob potential structure in DIII-D H-mode-like edge is not dipolar.

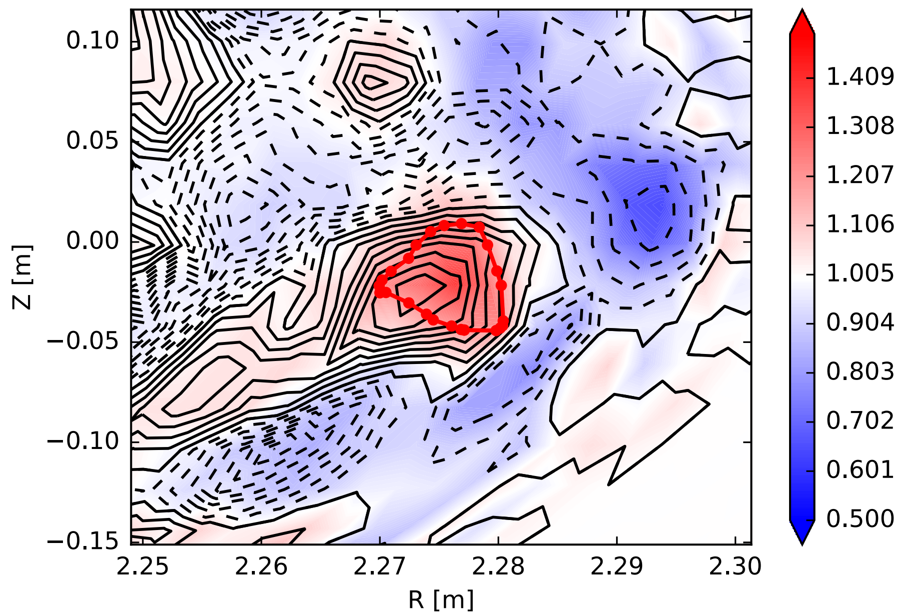

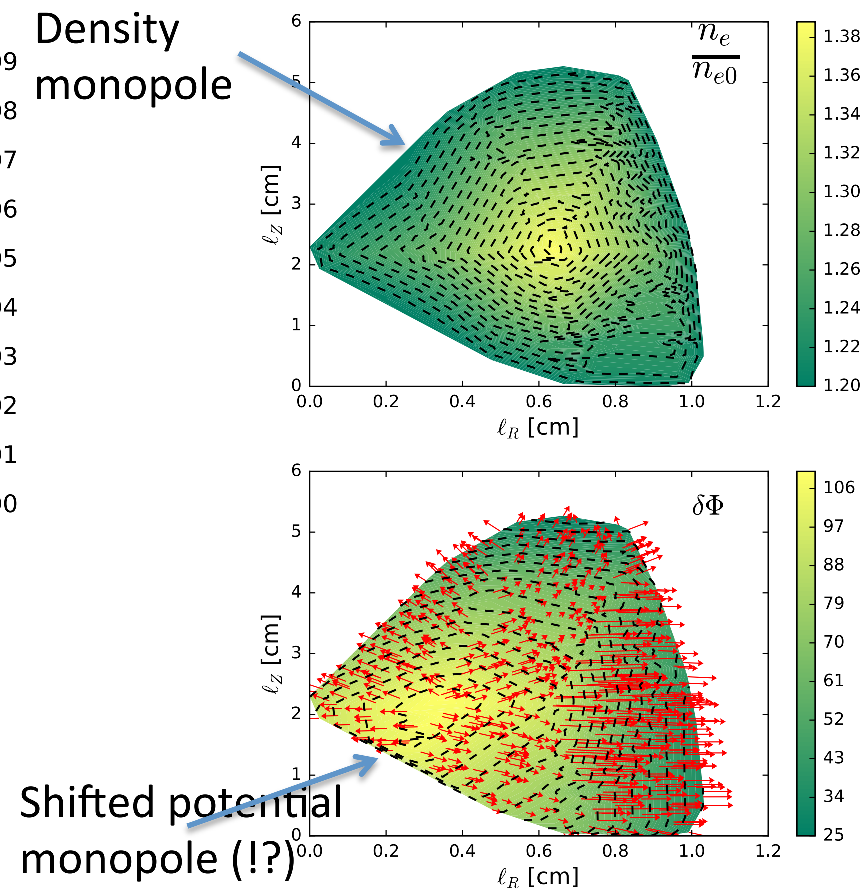

- Blob dipolar potential structure crucial to analytical blob models [Krasheninnikov PRL 2001, D'Ippolito PoP 2011]

[R.M. Churchill, *PPCF,* submitted]

Skewness and kurtosis of blobby δn/n

- Skewness and kurtosis of $\delta n_e/n_e$ increase near the pedestal foot, and into the near SOL
- Skewness slightly negative near pedestal top, similar to BES observations on DIII-D [Yan PoP 2011]
- High skewness in SOL similar to experimental observations with probes and BES [Boedo PoP 2003].
  - In XGC1, skewness begins increasing from pedestal foot.
- Non-Gaussian turbulence is related to L_n.

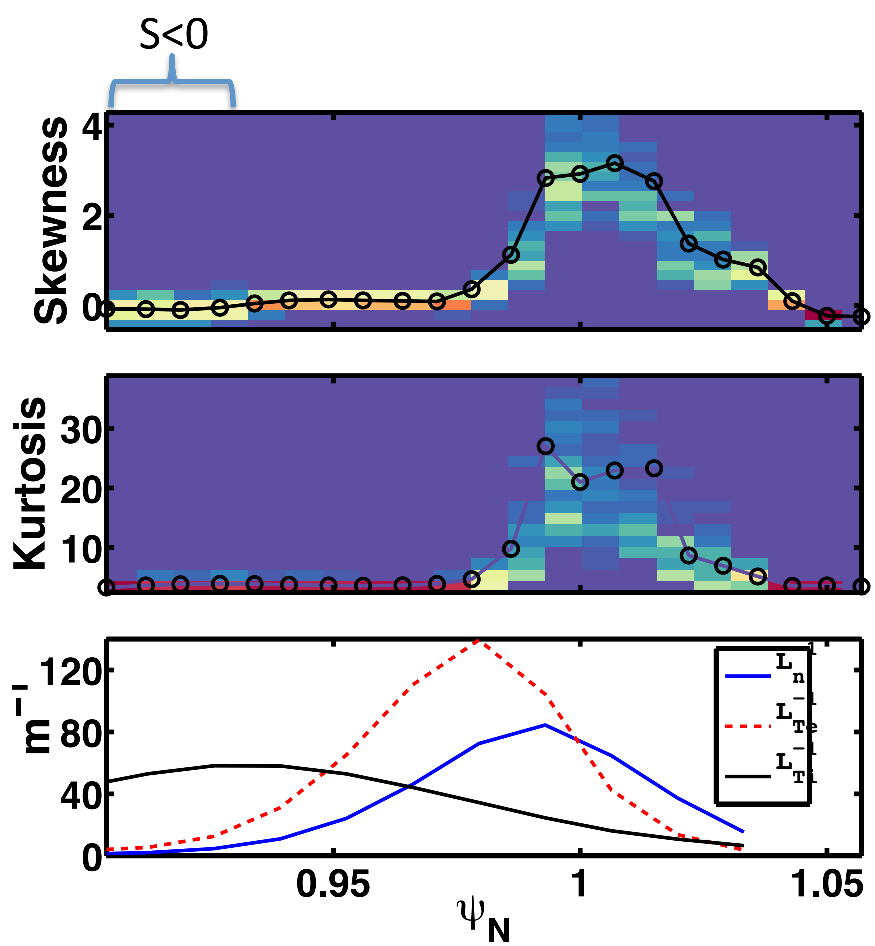

Blob's travel path is related to ExB

Blobs are born in the steep pedestal, cross the separatrix often with "shearing" action above the outer midplane, travel radially into SOL, and increase in relative intensity as convected radially, then poloidally to

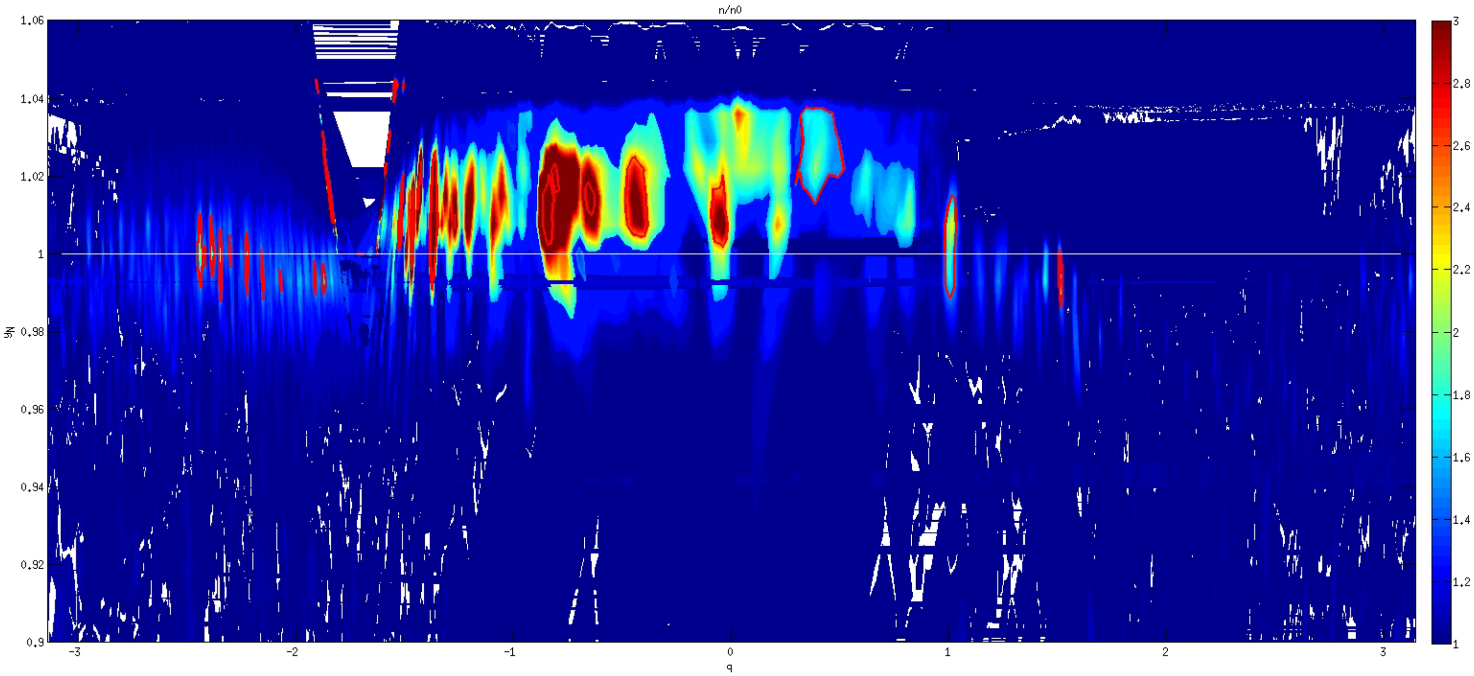

Frequency spectrum shows dominant turbulence-drive changing through the pedestal/SOL

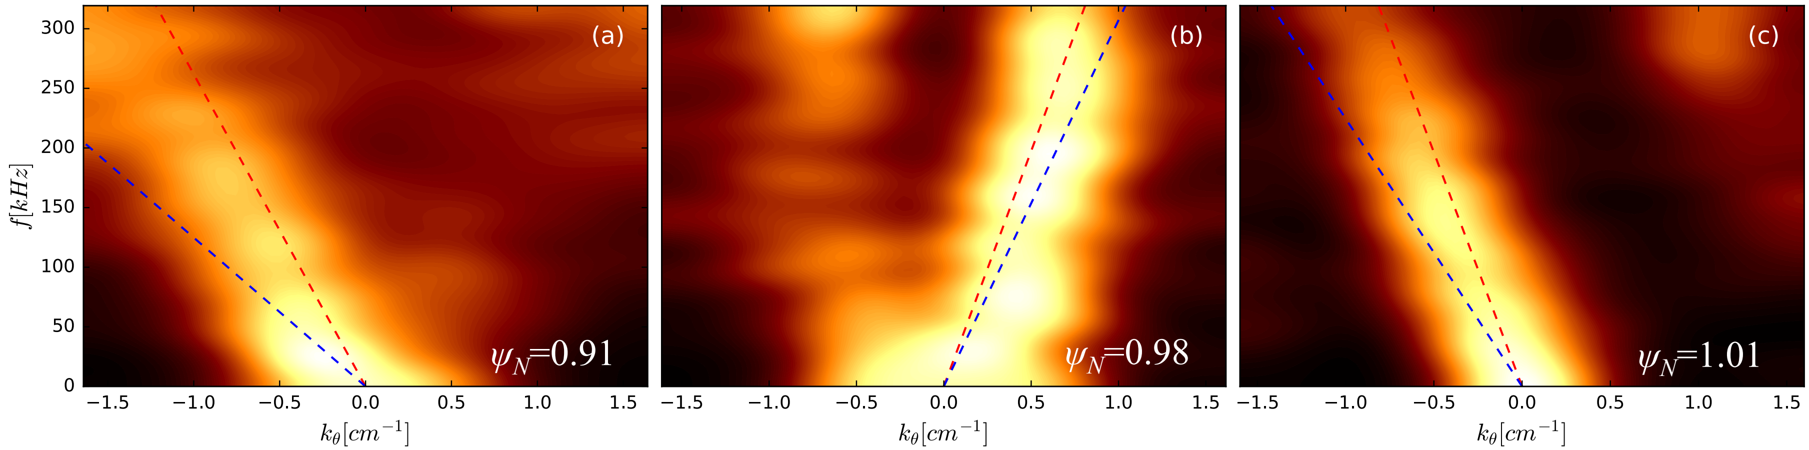

- Conditional spectrum S(k_θ,f) suggest dominant turbulence modes:
    - ITG near pedestal top
    - TEM through pedestal
    - Kelvin-Helmholtz type into SOL [W. Wang 2015]
- Dual propagating mode in pedestal region: nonlocal, counter-propagating turbulent structures [I. Cziegler, PhD thesis, 2012]

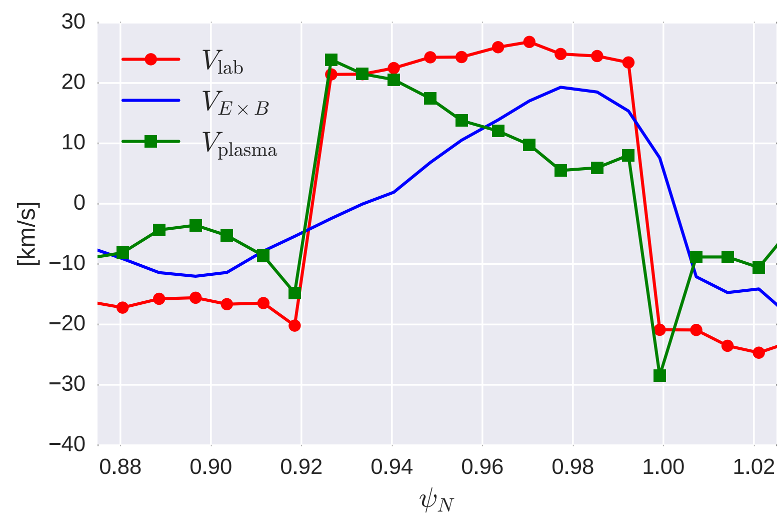

Myths about the neoclassical X-loss/X-transport physics

Myth #1: Is the X-transport theory the same as the previous orbit loss theories?

No. The previous orbit loss theories assumed that there is a large empty hole. In the X-transport theory, there is an unconventional transport process that closes the non-ambipolar v-space hole by ExB rotation and makes the collisional (+ turbulent) transport ambipolar [Chang, Phys. Plasmas 2002]

Myth #2: Is there strong momentum source from the X-loss?

No, only a little. The X-loss energy is raised so that the original v-space hole that can contribute to the momentum loss is closed [Seo, Phyis. Plasmas 2014]

→ Turbulence is needed to spread *f* into the higher energy loss hole.

Particle data are important

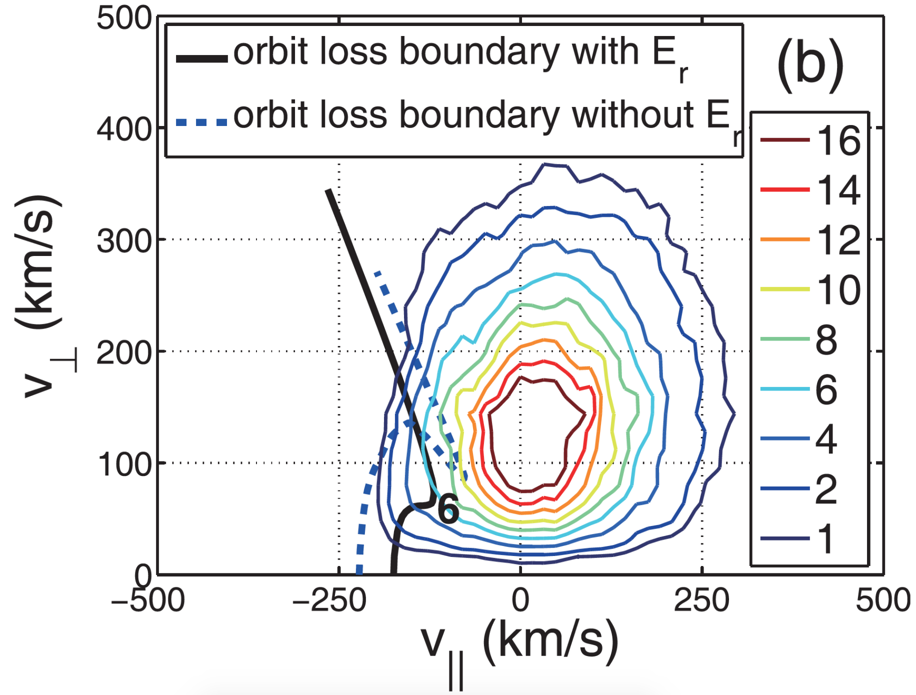

ITER-XGC1 λ_q in the Eich Chart

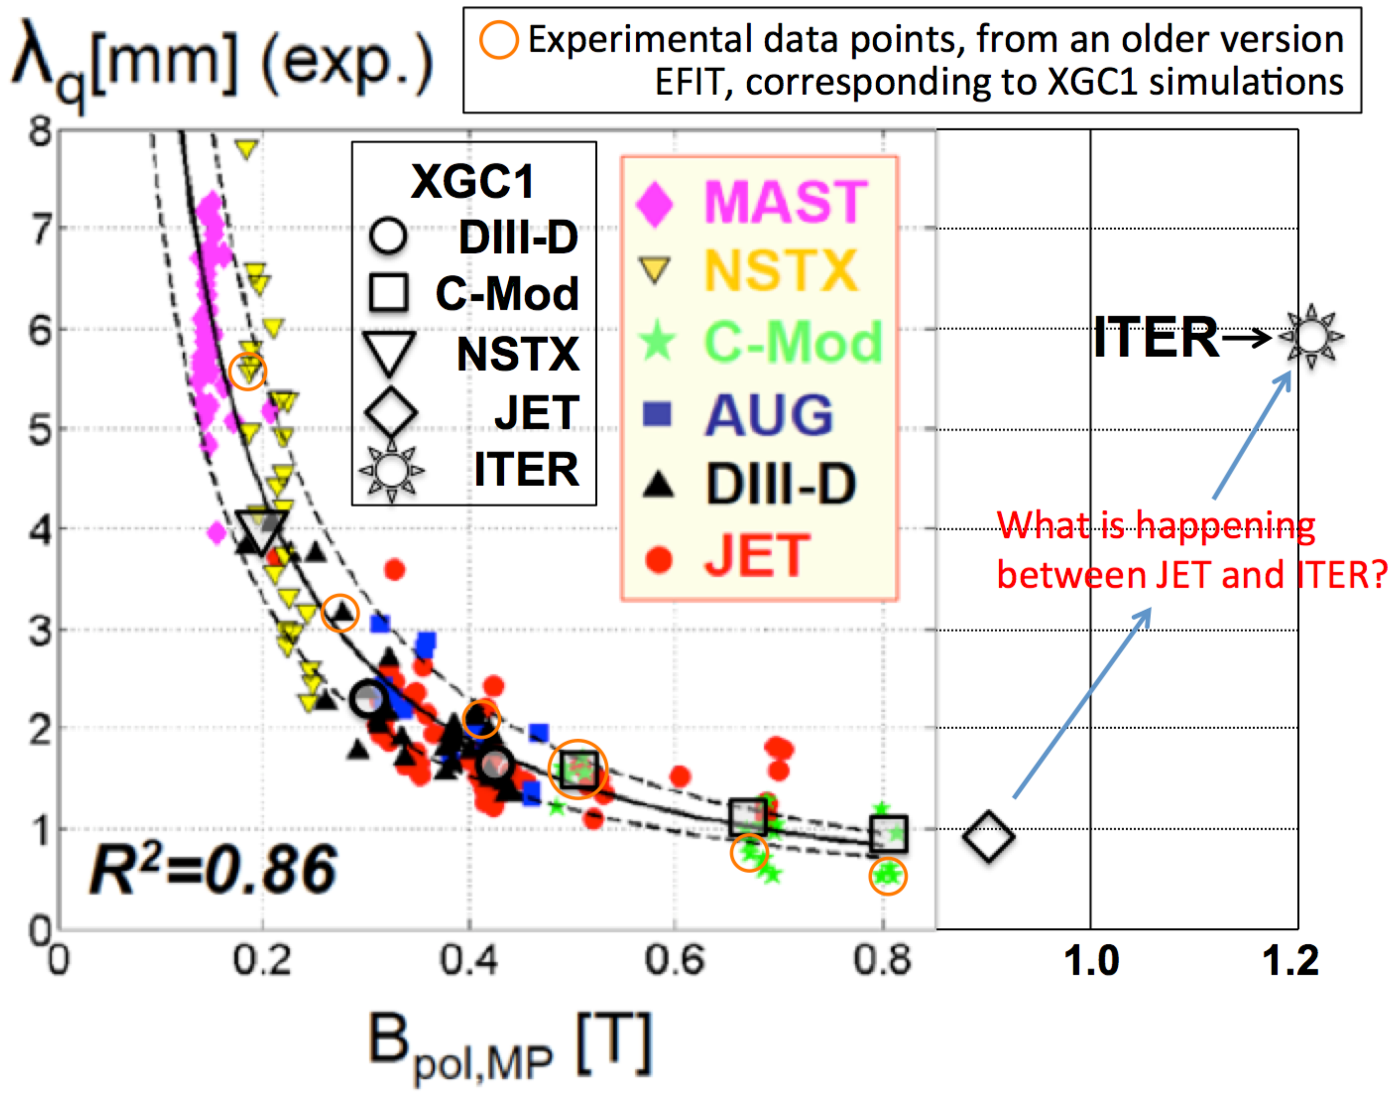

λ_q^{XGC} from the JET 4.5MA discharge follows λ_q^{Eich} within error bar

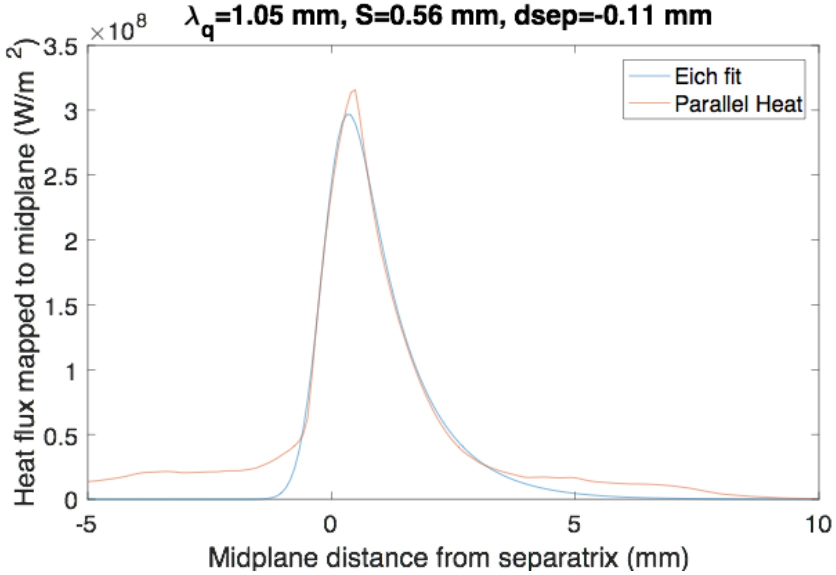

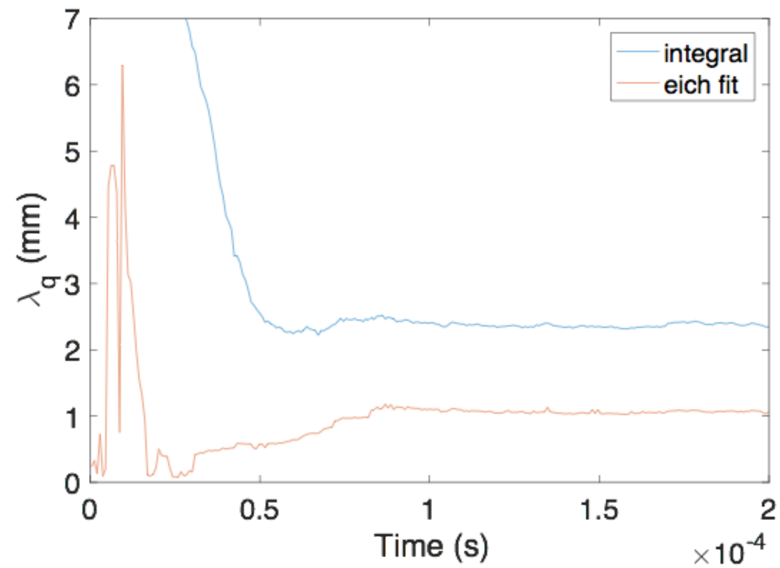

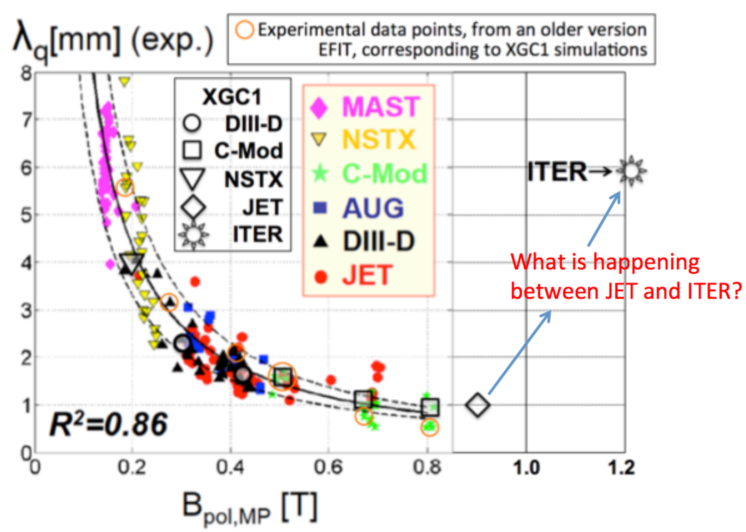

Physics bifurcation between present tokamaks and ITER

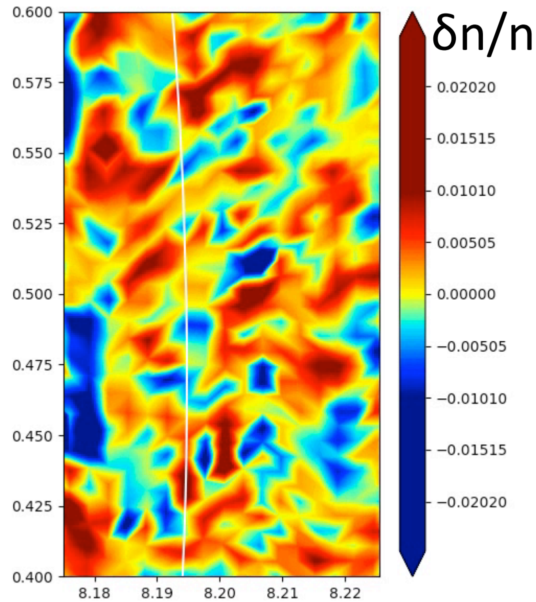

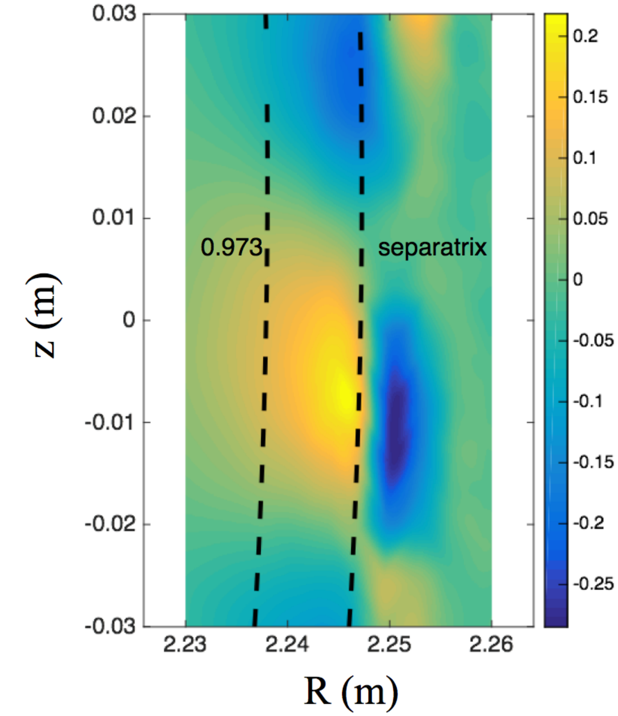

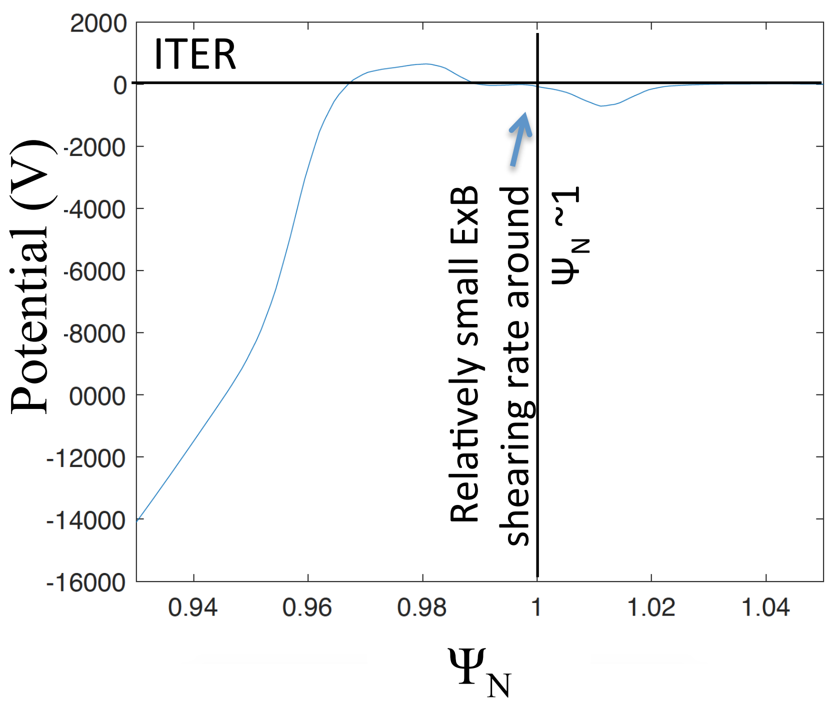

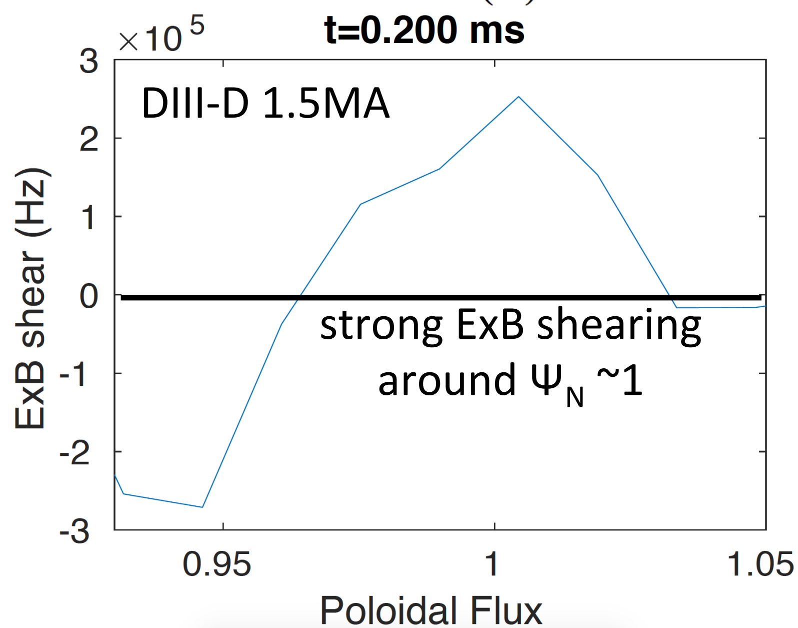

Ion particle-loss in ITER from inside the separatrix surface (Titan, 90% capability)

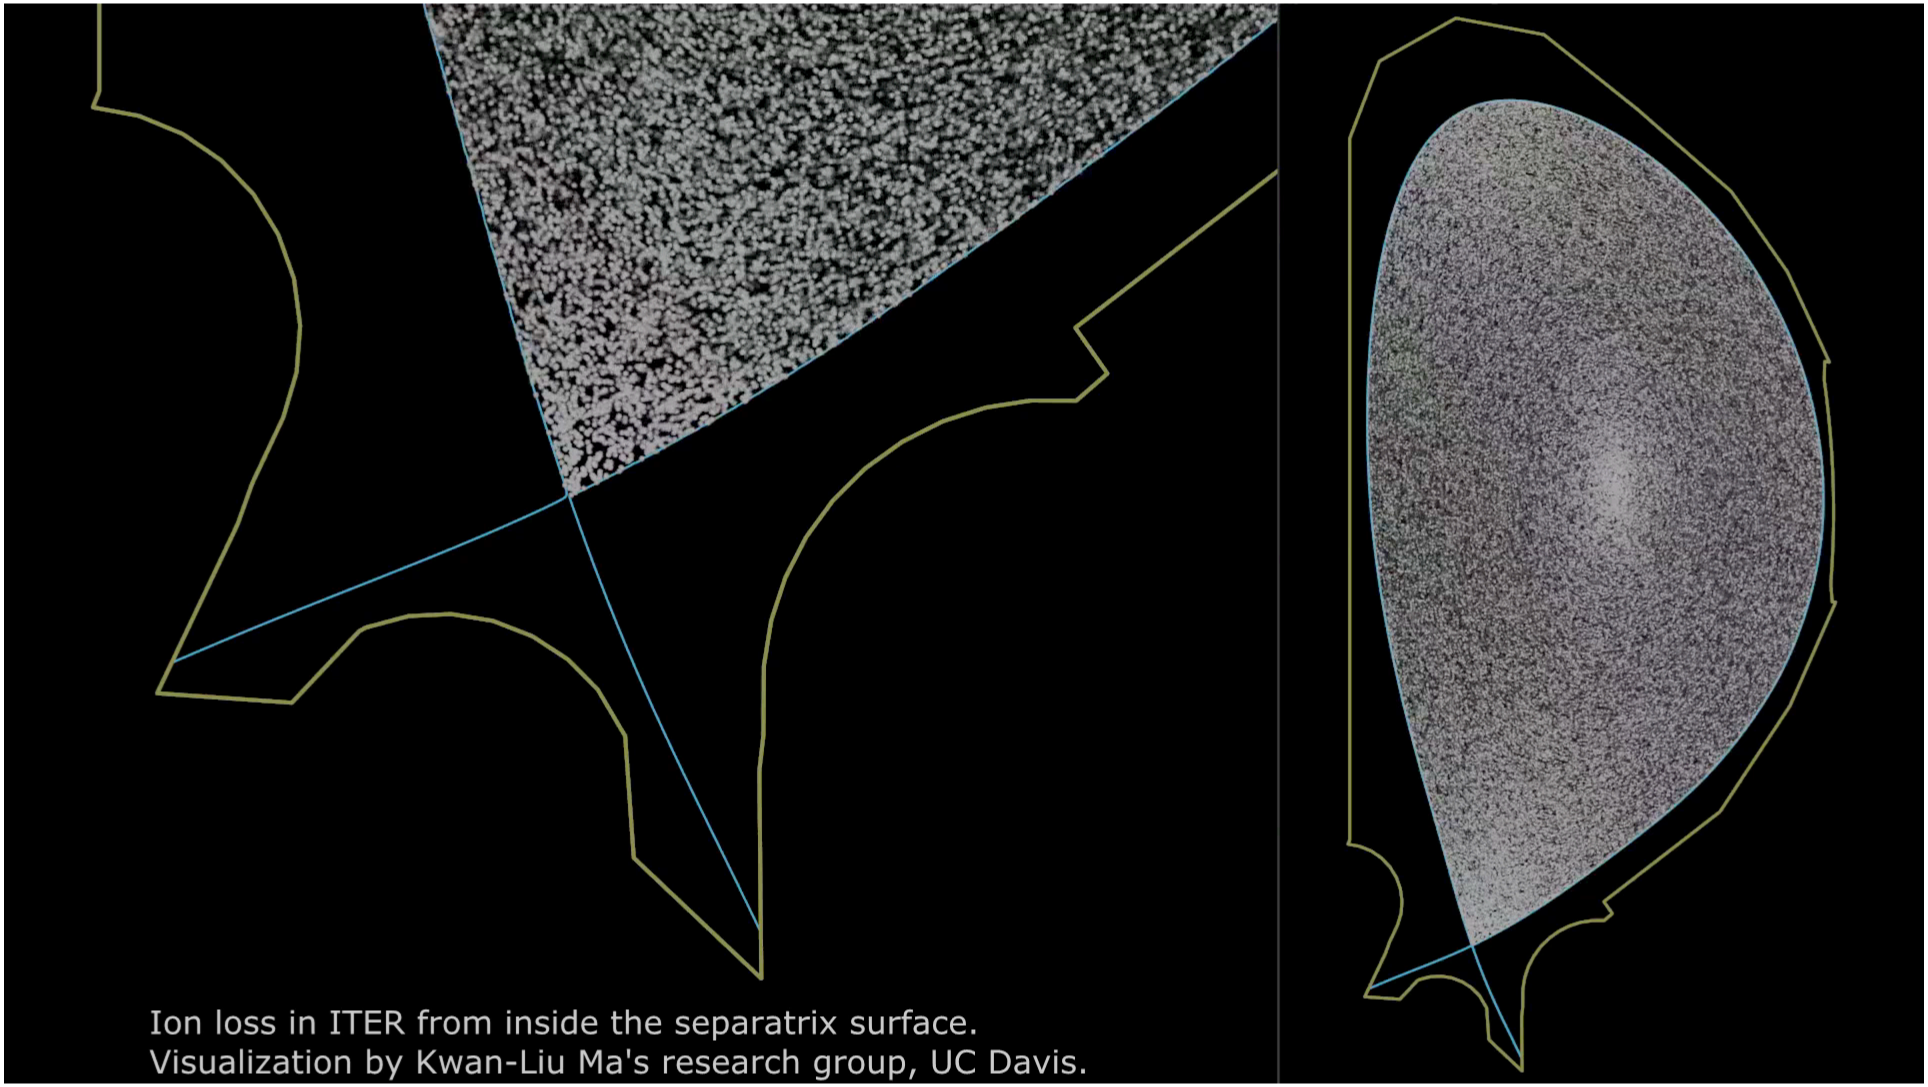

Gyrokinetic observation of L-H bifurcation in a C-Mod model plasma

1. t~0.175-0.21ms, suppression of lower frequency, higher amplitude turbulence occurs; and higher frequency, lower amplitude turbulence is generated (shades of green, eddy tearing by ExB shearing to be shown).
2. t>0.21ms, suppression of the lower amplitude turbulence follows.

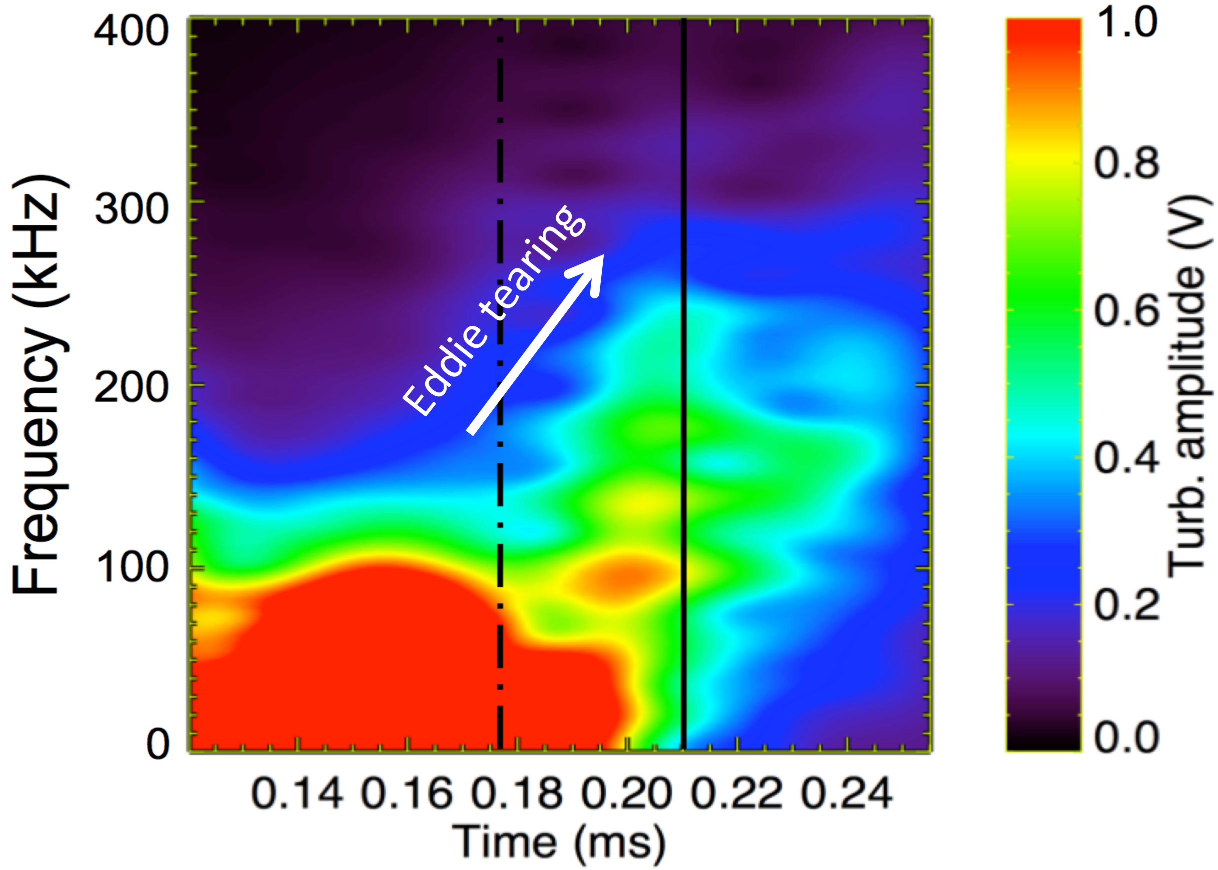

Forces that exist at the time of turbulence bifurcation

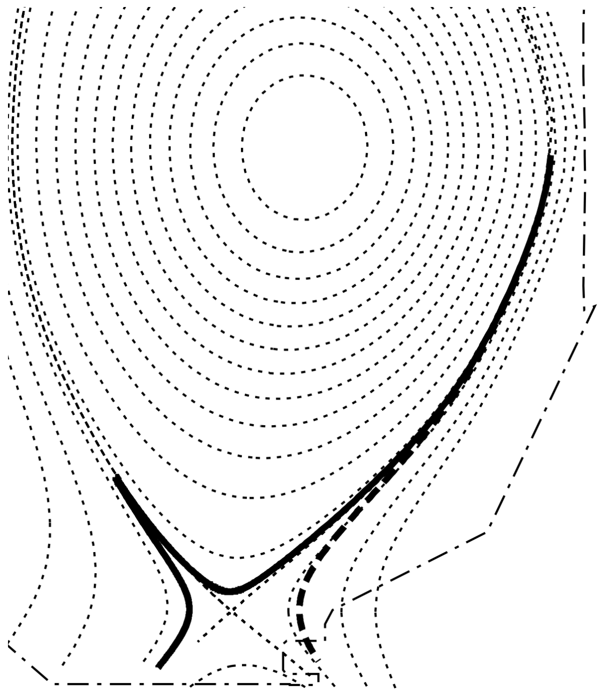

[S. Ku et al., PoP 2004]

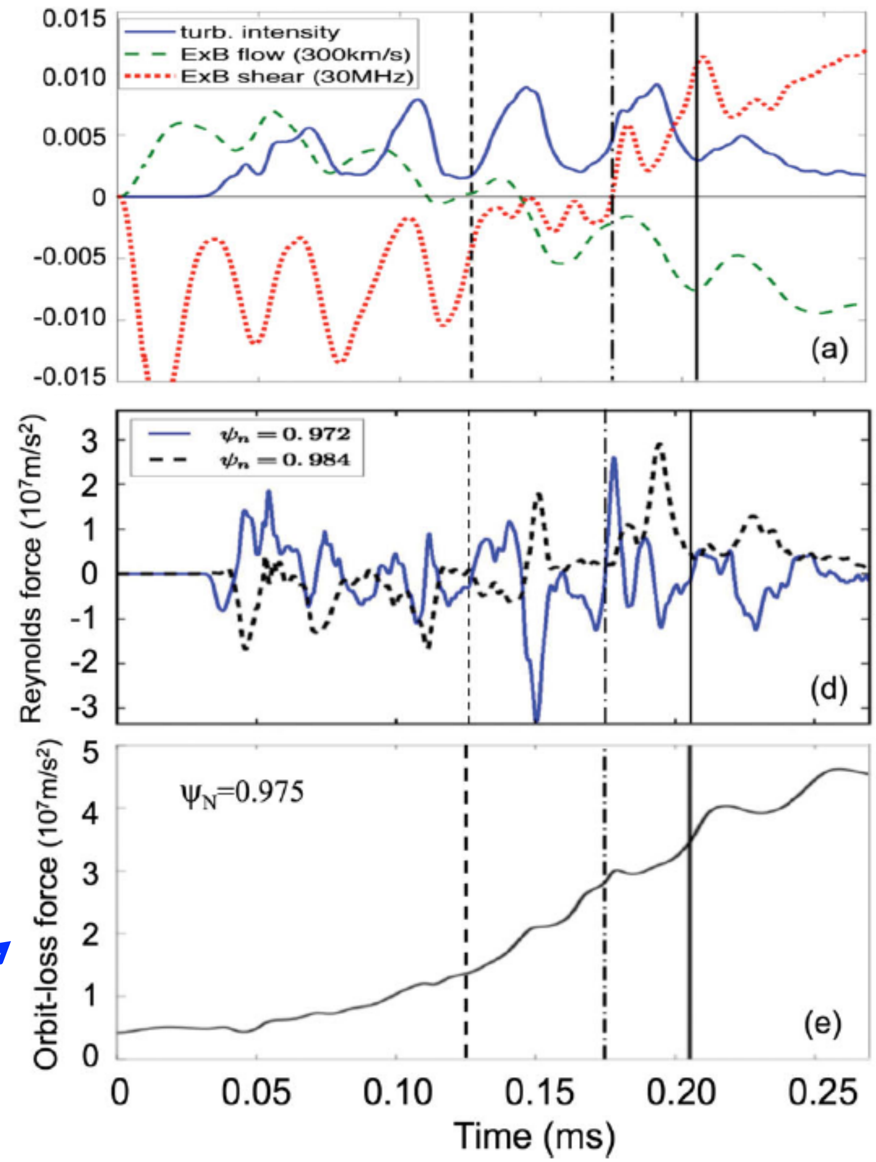

The orbit loss physics provides answers to all three questions.
[Chang, PoP 2002]

Strong anti-correlation (high density transport) between δn and δΦ in the near-separatrix/near-SOL region.

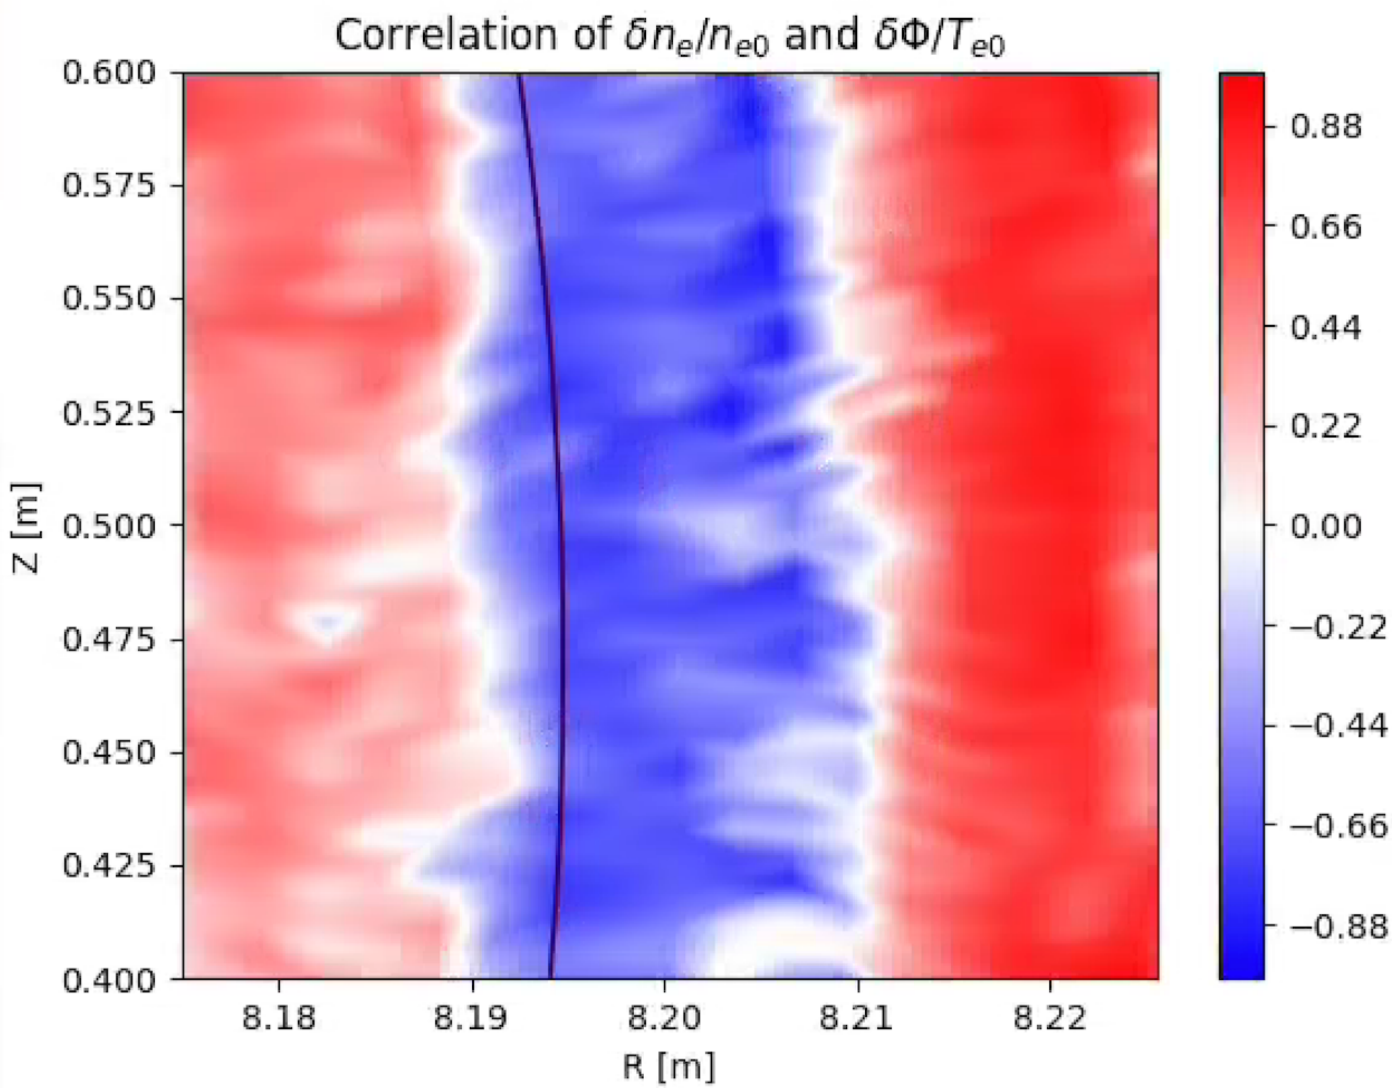

In ITER edge across separatrix, i- and e-turbulences co-exist

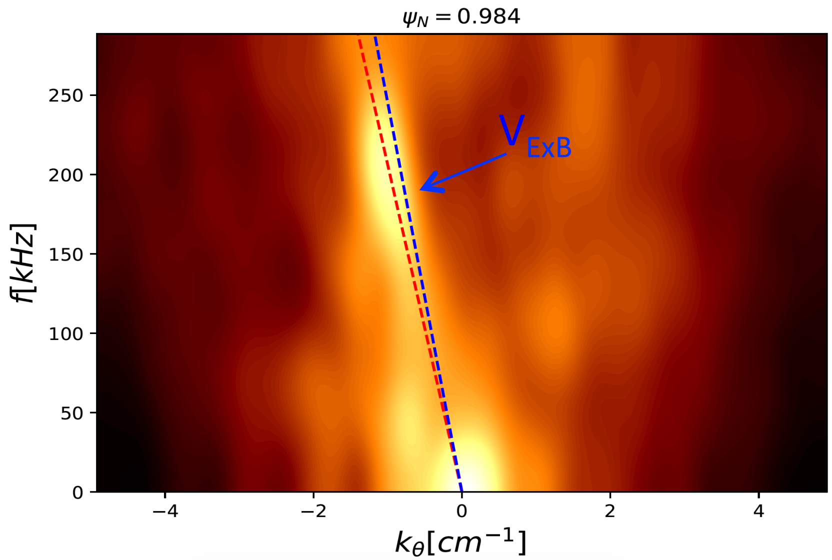

Linear study with GENE in the steep pedestal: If we remove trapped electrons, then the electron modes disappear.

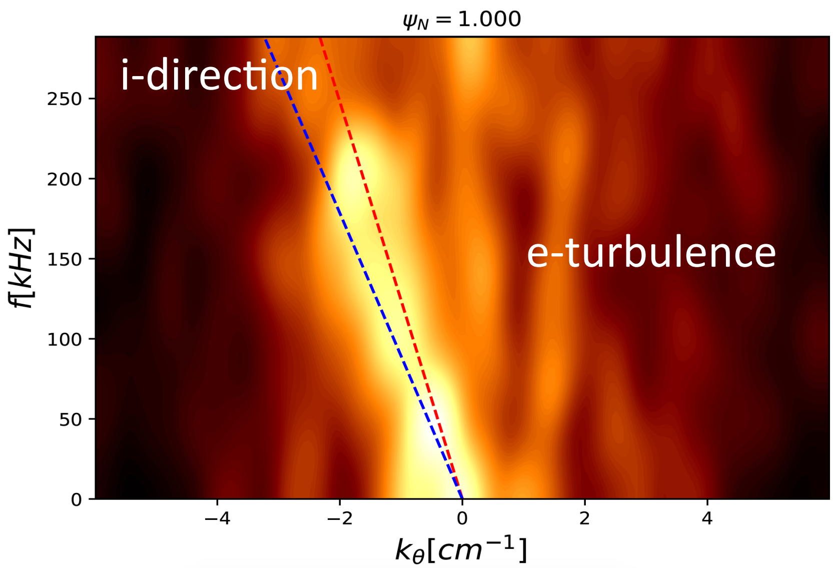

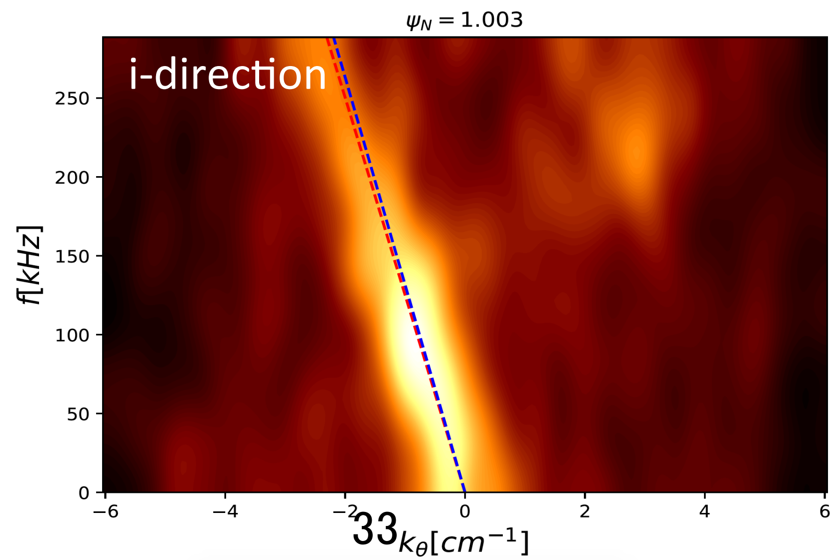

Data hierarchy in XGC1

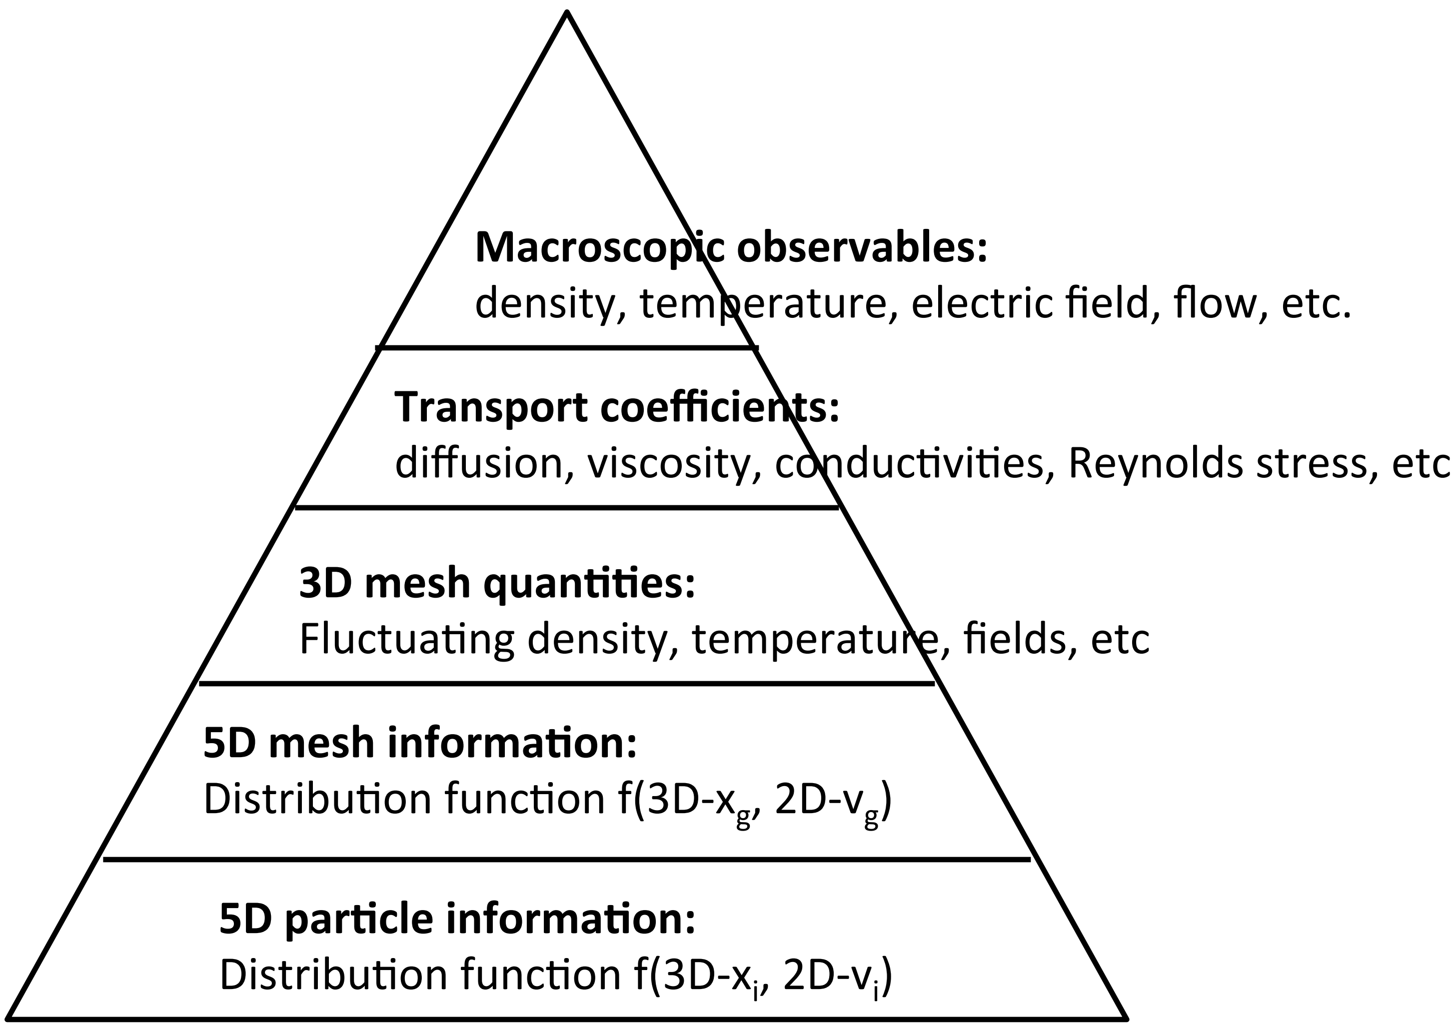

Big Data for today's ITER study in XGC1

For a reasonable ITER simulation today, with 10,000 particles per cell,

- 640B particles on 64M grid cells
- Total memory ~100TB at each time step = one check-point size
  - ✧ If 16,383 nodes are used (Titan) → 6GB per node
    - XGC1 is a low memory/node code
  - ✧ At 100GB/s using Adios, it takes ~17 minutes for one check-point output
    - Infrequent check-pointing
  - ✧ At the same time, an efficient management of particle and grid data on CPU and GPU memory is needed, given the latency
- First-principles physics study: Ideally, we would like to have all of the data to be written out (to avoid sampling error)
  - ✧ >1,000 timesteps for 24Hr run → >100PB per day >> 30PB, Atlas1+2
  - ✧ This amount data is unrealistically too big in all aspects
    - Write-out speed (>10TB/s >1TB/s Atlas1+2), size of temporary file system, movement to permanent storage, and data mining from storage
- Moving forward, we have no choice but to utilize
  - ✧ aggressive on-memory data analysis and reduction technologies
  - ✧ and NVRAM check-pointing, **on-the-fly**

A strategy used in the present XGC1 runs

- Write out only the coarse-grained f-data to file system

$$f = f_a + f_g + f_p$$

- ~40X40 velocity grids

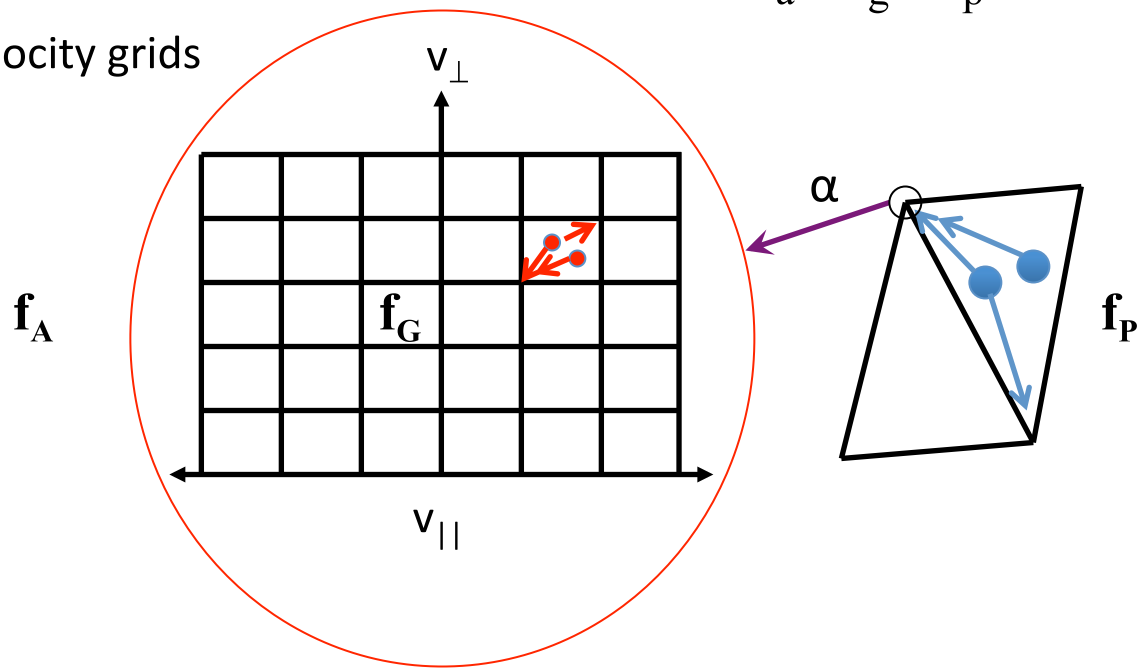

- Trying to analyze the fine-grained physics (from particles) on-memory, in-situ
  - Output the analysis results only
- Write out a small amount of down-sampled particle data for limited understanding of particle dynamics
- We lose most of the particle dynamics info.

36

Example for necessity of in-memory particle data analysis (Mira Highlight)

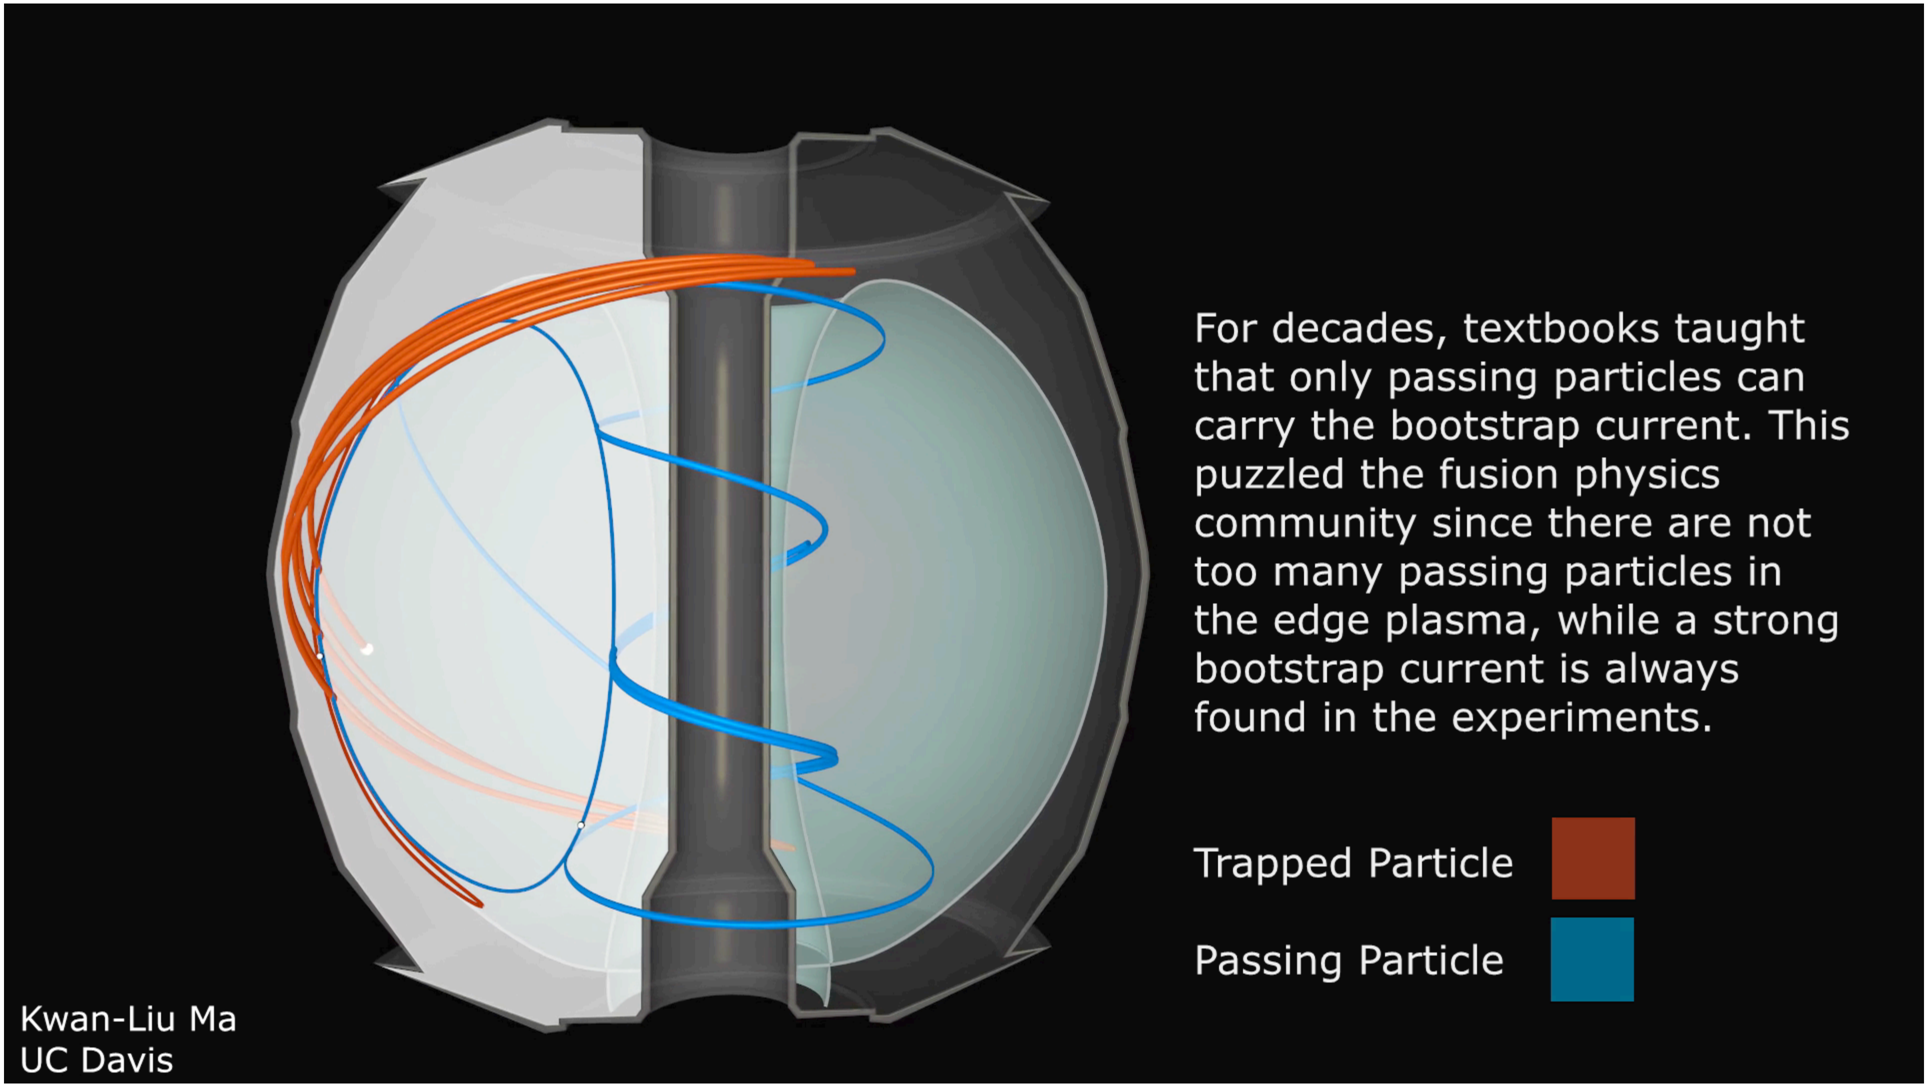

Strategy in moving forward

- Even with some creative strategies employed in the present XGC1 runs, the in-situ on-memory data analysis can produce far less physics analysis results than desired
  - Too costly since the main computing needs to be halted
    → On-the-fly analysis is needed
  - If all the local grid-node analyses are to be outputed for post-processing of inter-mesh operations, the output speed becomes too slow and the data accumulation becomes excessive in the file system → Use data analysis nodes with hybrid staging
- Higher resolution simulations together with more physics in the future will produce even bigger data and demands more aggressive in-situ data analysis/reduction technology

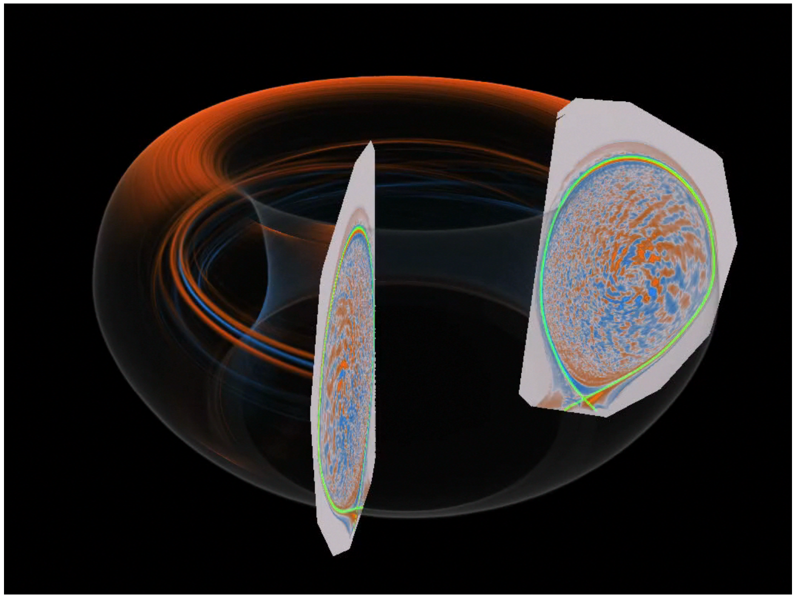

Reducing Payload Size

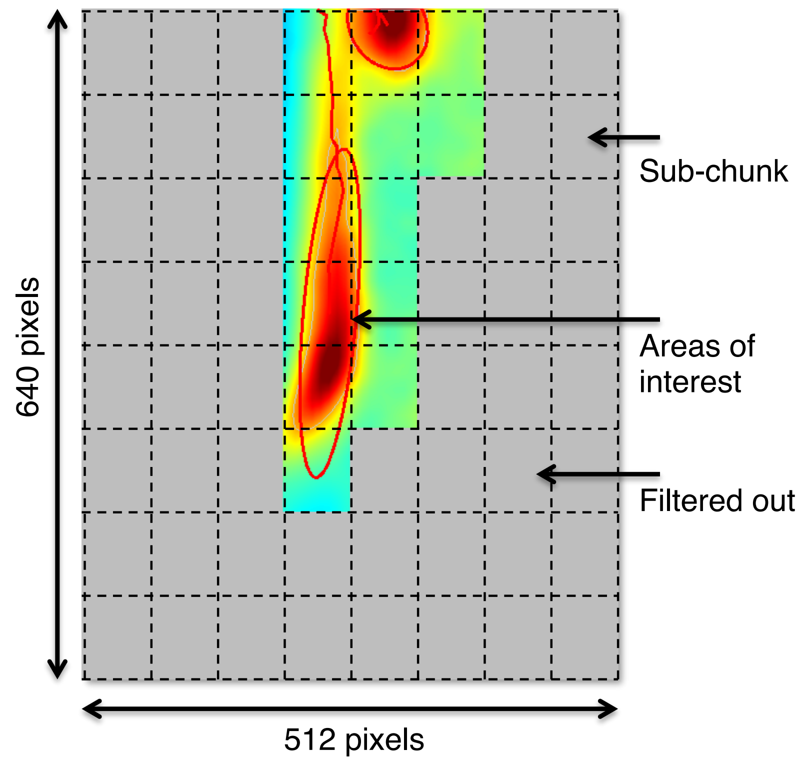

- Select only areas of interest and send (e.g., blobs)
- Reduce payload on average by about 5X

A new paradigm is needed in the XGC study:

Runing XGC will be like running a large international experiment.

- International community planning
- Centralized execution on extreme scale HPCs
- International community analysis of data (real time)

Difference is the uncontrollably Large daily rate/amount of data.

Big Data from ITER experiments

- **Volume**: Initially 90 TB per day, 18 PB per year
  - maturing to 2.2 PB per day, 440 PB per year
- **Value**: All data are taken from expensive instruments for valuable reasons.
- **Velocity**: Peak 50 GB/s, with near real-time analysis needs
- **Variety**: ~100 different types of instruments and sensors, numbering in the thousands, producing interdependent data in various formats
- **Veracity**: The quality of the data can vary greatly depending upon the instruments and sensors.

Difference from experiments in other areas: Near-real time feedback for "daily" steering

The pre-ITER superconducting fusion experiments outside of US will also produce increasingly bigger data (KSTAR, EAST, Wendelstein 7-X, and JT60-SU).

Data generation rate can be exceeded by first-principles simulations

Computer Science Technology can accelerate ITER Science

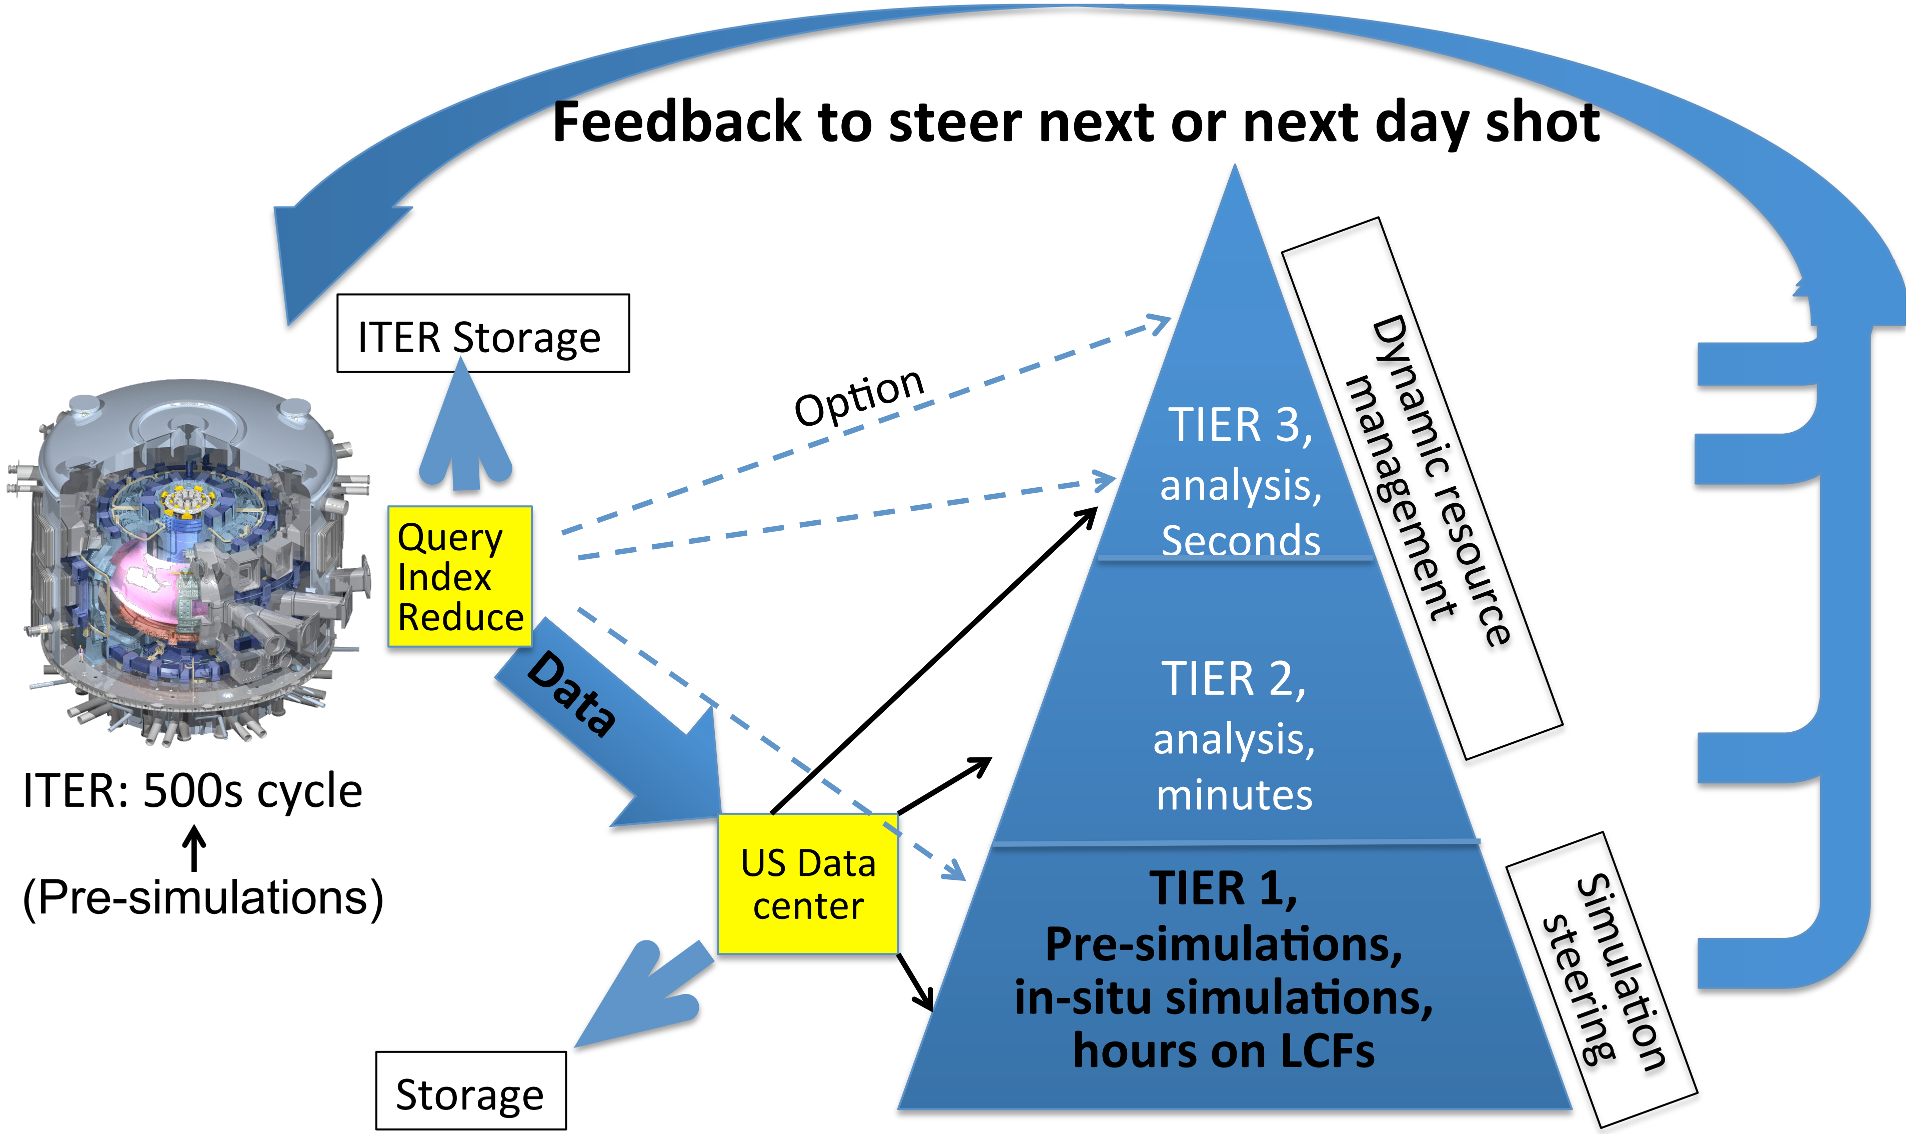

Faster success of ITER program and earlier development of commercial fusion reactors

Example ITER workflow:
Anomaly detection, analysis, and feedback

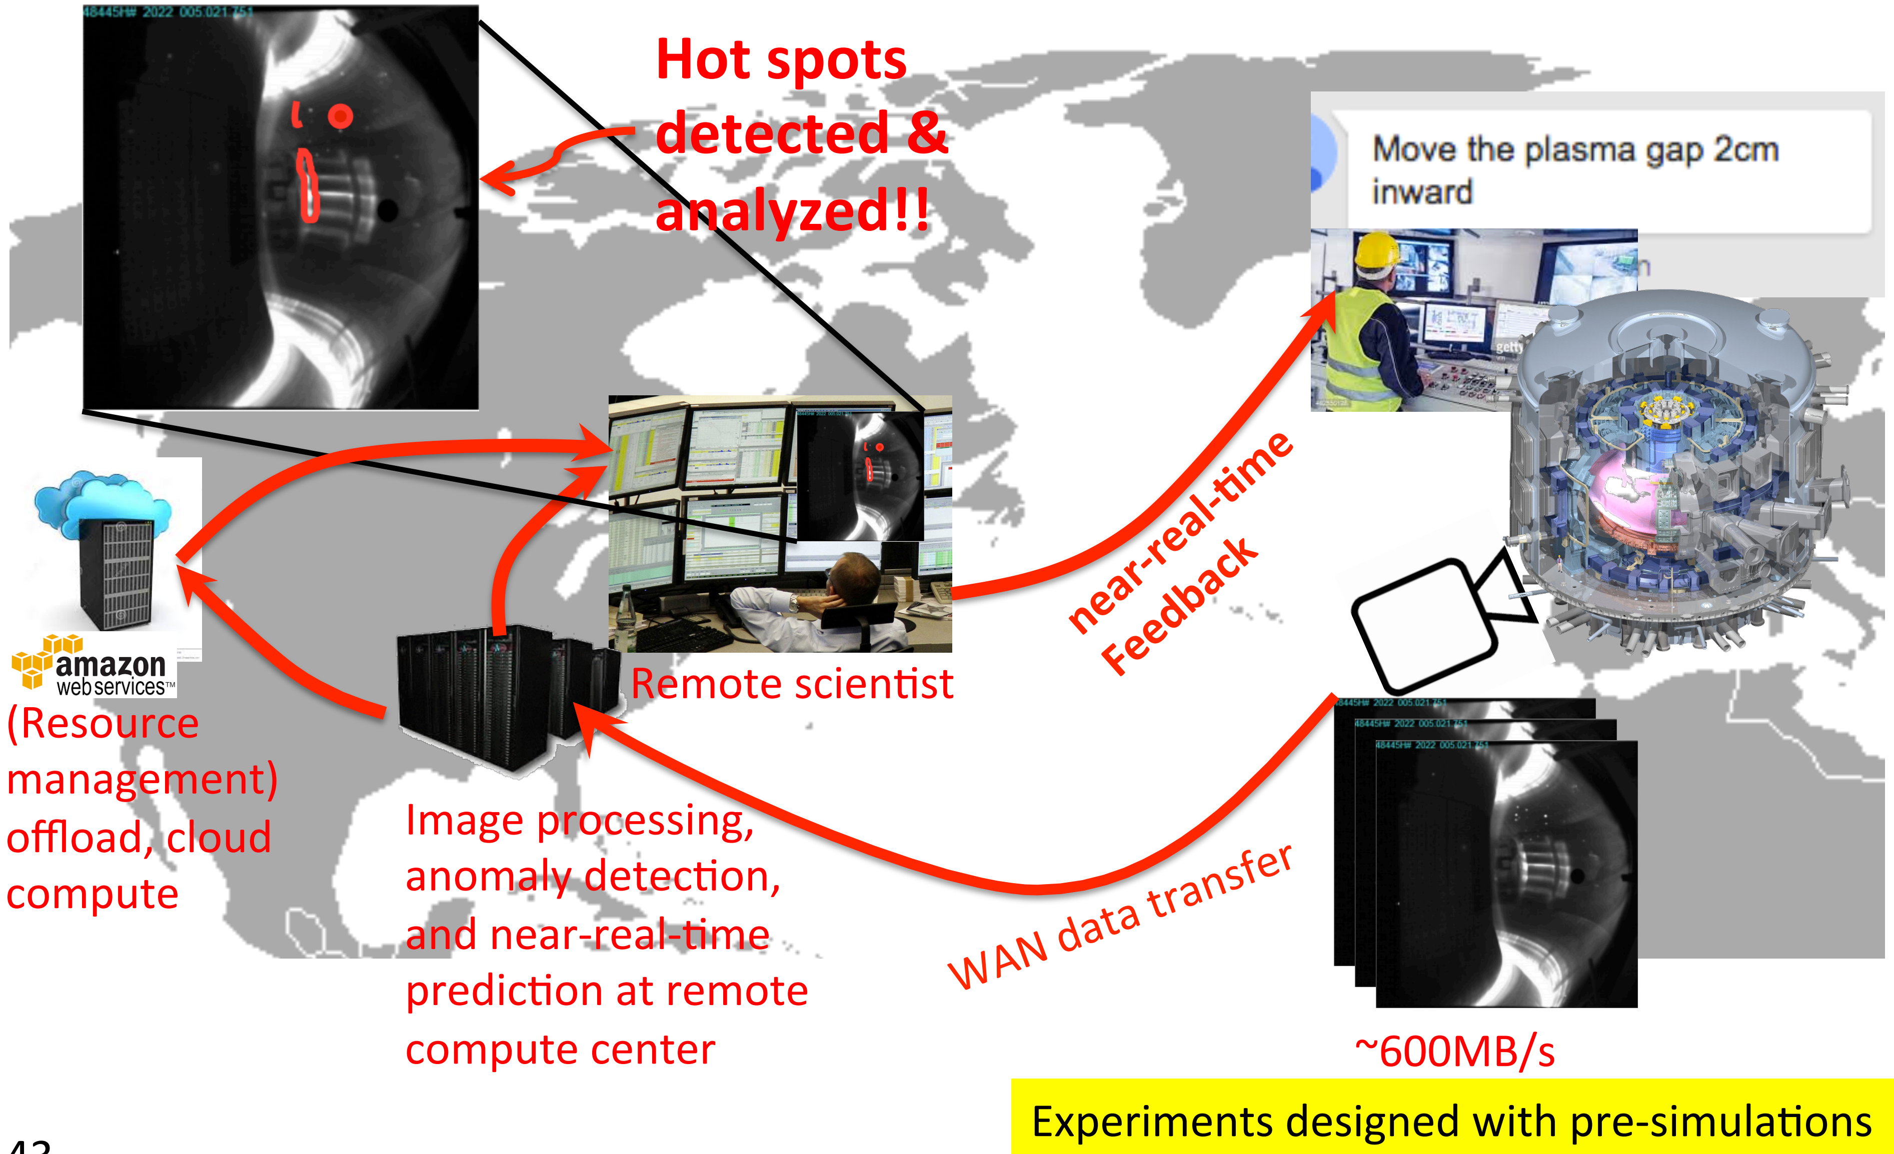

Summary

- XGC data is too big for post-processing
- Our mission is to discover new physics, and use the present knowledge to look for them. But, eveutually, we may run into discoveries that we may even not know what we have to look for. → Machine learning.
- We need help in data reduction, analysis, feature detection, visualization, etc.

→ Physics, math, and computer science need to work together.